David Hawkins Research

Forensic Economics
Mentoring Communities
Virtual Prototyping and A.I.

DCM-001: This paper develops a theory as to how experts think, acquire, understand, and use their knowledge (particularly with reference to petroleum geology). The theory identifies a number of human, knowledge-handling techniques which have been successfully implemented in the Deductive Computing Machines (DCM) — invented by David Hawkins — and adapted to meet user specifications from any **"real-world"** knowledge base.

"AN ANALYSIS OF EXPERT THINKING"

by David C. Hawkins

Sclumberger-Doll Research
Ridgefield, Connecticut (U.S.A)

"This note is simply to say how much I enjoyed your paper in IJMMS, "An Analysis of Expert Thinking". It was a refreshingly vibrant approach to a subject that Artificial Inteligence "knowledge engineers" have virtually castrated. It's good to see expertise baring its fangs once again, so to speak."
- Dr. B. Webber, University of Pennsylvania, Moore School of Electrical Engineering

Published in the International Journal of Man-Machine Studies (1983)18, 1-47 London:Academic Press

For further information or to invite David Hawkins for a speech or workshop on the Analysis of Expert Thinking, please contact:

Email: DCM-001@DavidHawkinsResearch.com
Website: www.DavidHawkinsResearch.com

Mail: David Hawkins Research, DCM-001
c/o #202- 2001 East 36th Ave.,
Vancouver, B.C., Canada V5P 1C9

AN ANALYSIS OF EXPERT THINKING

by

David Hawkins

Schlumberger- Doll Research, Ridgefield, Connecticut (U.S.A.)

ABSTRACT

Human expertise should be better understood before the users of expert systems specify the services needed and expected from such systems. An analysis of expert thinking has been developed to assist in this understanding. The analysis is discussed in the paper under three main headings:

Specifications: examples are given of the services users obtain from human experts, in the particular domain of petroleum geology. These services indicate general qualities desirable in a human and, by analogy, in a system. The qualities are listed as specifications for expert system design.

A Theory of Expert Thinking: how human experts acquire, understand, and use their knowledge (particularly with reference to petroleum geology). The theory identifies a number of human, knowledge -handling techniques which could be implemented in a system to meet most of the user's specifications.

Human and System Expertise: a comparison suggests that, properly designed and suitably applied, an expert system can help its users make well- informed decisions; failing this, the system may prove dangerously misleading and should not be accepted as a substitute for an accountable, human expert.

TABLE OF CONTENTS

An Analysis of Expert Thinking

SECTION 1: INTRODUCTION

1.1 INTRODUCTION

This paper has been written from the personal viewpoint of a petroleum scientist who is a potential user of human or system expertise in the field of petroleum geology. My present interest in expert systems comes after a number of years working in the hydrocarbon exploration industry. This work brought me into contact with many professionals and specialists and, occasionally, with people who classed themselves and were recognized as experts. To communicate with such experts was an impressive experience and some of the skills that distinguished them from those who used their expertise were readily apparent.

A number of expert systems have developed from recent Artificial Intelligence research into theories of the mind. The mental activities and services performed by the human experts of my experience, have suggested a particular theory described in this paper. I am not, however, attempting a direct criticism of the principles of these established systems. To a potential user, it proved more productive to select a familiar domain and examine the make-up of the associated expertise. It was then possible to speculate on how humans provide their services and the extent to which systems could do the same. Given the results of this study, other users should be in a better informed positions to accept, modify or reject the expert systems offered them.

Major benefits accrue to the cautious users of the best human expertise: significant risks are run by those who put much faith in a mechanistic diagnosis of a complex problem in nature. I have, therefore, proposed a set of system specification that should be met before users agree to switch loyalties from their human colleagues. Subsequently in the paper, I develop a theory of expert thinking which suggests how such systems might be implemented. Finally, the theory is used to make a comparison between the human’s and the system’s ability to “understand” the human users’ problems.

In the abstract, the benefits and dangers of expert systems are not self-evident: petroleum geology is a convenient, applied science for a more graphic illustration. The domain is used throughout the paper to demonstrate the desirable and, perhaps, essential characteristics of any expert system aimed at maximizing the benefits to its users and minimizing the risks.

Petroleum geology is a domain that would benefit from the application of an expert system, irrespective of its value in illustrating various ideas on expert thinking. Before launching into the analysis, it is worth considering why.

1.2 PETROLEUM GEOLOGY: PROBLEMS OF COMMUNICATIONS AND KNOWLEDGE

The proper analysis of information in the field of petroleum geology is hindered by the degree of specialization required from the analysts. The number of scientific and professional disciplines involved in the location and evaluation of the world's hydrocarbon reserves, illustration the problem. A conservative list of these disciplines would include:

Geology	Geophysics	Seismology	Mineralogy
Log analysis	Sedimentology	Stratigraphy	Paleontology
Geochemistry	Economics	Statistics	

Concerned personnel are almost obliged to be specialists in the above disciplines if they are to handle the sheer volume of information generated by modern oilfield data-acquisition techniques. As specialists, they must still communicate with each other because the geological objects of their interest exist in nature: a nature that operates as a complex whole; a nature that behaves independently of any divisions made to it for the sake of human and academic convenience.

Companies and institution frequently combine two techniques in an effort to encourage better communications amongst their specialists and so improve the quality of their geological interpretations. One effort is to broaden the base of specialized knowledge within the organization through practices like job rotation, more training, or the creation of multi-disciplinary project teams. Alternatively, or in addition to this effort, an organization may employ experts to review the various interpretations offered by its specialists and construct acceptable composites. Most specialists would agree that the first approach meets with limited success: the time they require to deepen their knowledge of a subject seems to take precedence over the time, and inclination, needed to understand the role that subject plays in the broader issue. The use of knowledgeable experts, negotiating agree interpretations with the specialists, is a better solution. It is only a practical solution, however, if these experts can be found and as long as the range and complexity of the knowledge they require, remains manageable.

Petroleum geology has long passed the stage where experts can hope to keep up with the "explosion" of knowledge in the contributory sciences. There have

been painful and expensive consequences for the oil-field industry. A shortage of experts has lead to a concentration of the expertise that does exist into areas of limited, regional application. Research scientists rarely have the chance to discuss the practical consequences of their work with adequately-qualified people. Professional users are falling further and further behind in applications of the latest analytical techniques.

Petroleum geology is therefore an appropriate subject for the application of an expert system. Extensive and detailed knowledge of the domain is required to make a "best" interpretation of frequently ambiguous, geological information. Human expertise, although scarce, is available, so that the way in which it is used can be studied. Petroleum geology offers the expert repetitive but subtly distinctive problems. Any system built, can be monitored to see whether, like a human expert, it can improve its interpretation techniques by experience, or whether it is too rigid to handle subtle requirements and so compare the quality of the service from the system with that provided by the human.

1.3 ORGANIZATION OF THE PAPER

The three main themes of this paper are presented in Sections 2 through 4.

Section 2 describes some examples of the services expert petroleum geologists appear to provide to their "clients". These examples point to the qualifications such clients might reasonable request, if they were trying to recruit a human expert. The qualifications have been translated into a set of specifications which an expert system would have to satisfy, in order to provide a quality of services similar to the human expert.

Section 3 develops a theory of expert thinking, via an analysis of the techniques humans use, albeit unconsciously, to learn and understand information, propagate effects, handle conflicts, and finally, to justify their results to the users' satisfaction.

Section 4 examines some different types of user interactions from the position of the expert, both human and system. This examination distinguishes areas where a system might succeed, and also where it would probably fail, in trying to understand its users' needs.

The remaining sections appear as follows: Section 5 summarizes my theory of expert thinking and concludes with the theory's implications for the design of

expert systems; Section 6 is a list of references; and, Section 7 is an acknowledgment to those who have assisted me in writing this paper.

1.4 TERMINOLOGY

A rather specialized vocabulary is required for a study of expert thinking, particularly when the hypothetical expert is thinking about problems in a specific domain such as petroleum geology. An attempt has been made to keep the jargon in this paper to a minimum. Unfamiliar terms are not defined if their meaning is fairly obvious from the context in which they are used. Otherwise they, and familiar terms used in a restricted sense, are furnished with bracketed explanations early is their appearance as text. These explanations are mostly selected, and where necessary, modified, definitions taken from the reference books listed at the end of this paper.

The word used in the title and frequently throughout the paper is used in a restricted sense, deserving a preliminary explanation. “Expert” has a dictionary definition of “someone who has acquired a special knowledge in a particular subject.” I understand and use the word to mean “someone who can negotiate an agreed interpretation of a particular subject with the help of special knowledge and user opinions.” With this definition, an expert appears very much as an analytical tool, helping the users make well-informed decisions without forcing them to accept any particular interpretations or procedure.

SECTION 2: SPECIFICATIONS

2.1 SETTING THE SCENE

The objective in this section is to take the reader through an exchange between an expert and users of expertise and so identify what it is that the latter obtain during such a transaction. The section concludes with suggested specifications for an expert system: these specifications would have to be met if the users were to obtain something approaching human-level services from the system.

A good way to study the perceptions users may have of such services is to imagine a threesome, furnished with the same information and trying to resolve a fairly realistic problem of descriptive geology. The trio is made up of an expert, a geologist and a log analyst.

A distinction is helpful here as the first example of the exchange is limited to a problem of descriptive, as opposed to interpretative, geology. Initially, the group is only looking at answering questions of the “What?”, “Where?”, or “When?” variety. Such questions are essentially answered by describing the properties of the object of interest which are relevant to the question. The second example of this section illustrates the expert’s response to the “How?” or “Why?” type of questions. These latter questions require genetic (determined by the causal antecedents of something) explanations as answers.

2.2 EXPERT GEOLOGICAL DESCRIPTIONS

Imagine a well has been drilled through some underground, sedimentary formations in which a certain rock formation, called Berea sandstone, is believed to exist. “Berea,” as a name, indicates that the geographical location of the formation. “Sandstone” indicates the rock in this formation is mainly composed of quartz mineral (chemical formula SiO_2).

A typical geological description of the features of Berea sandstone might appears as follows:

COMPOSITION:	Mainly quartz (SiO_2) mineral. Less than 5% by volume of feldspar. Less than 20% volume of formation fluid.

GRAIN SIZE:	Fine to medium.
FABRIC:	Grain contact type of tangential or long.
LOCATION:	Berea.
OTHERS:	Nothing particularly associated with Berea.

Believing the well has somewhere intersected a Berea sandstone formation, the group still requires additional information before any description can be attempted of the intervals drilled. They are therefore provided with a geophysical log that is a record of the density of the formation along the well borehole. (Devices used for geophysical logging provide continuous records, against depth, of some of the physical properties of the rock being measured.) The formation density log, in this example, is a record of the rate of detection of gamma rays that have passed through the formation after emission from a source in the device. This rate of detection can be shown to be dependent on the density of the formation, hence the name of the log. Apart from the density log, the trio is also furnished with rock samples, called cuttings, that have been brought to the surface by the drilling process.

Suppose the expert is asked by one of the users, "Where – at what depth – is the Berea sandstone in this well?" Given the information available, the problem can be approached in two ways: the classificatory or the behavioral. (This distinction is discussed later in a more theoretical sense.) In practice, the classificatory approach requires that the cuttings, from known depths, be checked for features typically associated with that class of rocks, named Berea sandstone. The behavioral approach requires the density log to be checked for depth intervals which the apparent rock density behaves (varies with depth) in a way consistent with Berea sandstone. As explained below, the expert faces an additional complication: both the cuttings and the density log may not be representative of the properties of the rock formation from which they were, supposedly, taken or recorded.

The rock cuttings, brought tot the surface by the drilling mud, have been broken off the formation by the drilling bit. The drilling process may have significantly distorted the original grain size and fabric of the undisturbed rock. Depending on their density and shape, and the viscosity of the drilling mud, some sorting and mixing of the solid particles takes place as they move up the well with

different vertical velocity. The samples, collected at the surface and given to the geologist for examination, are picked off a sieve that inevitably allows some very-fine grained material to pass through. An estimate of the depth from which the sample came could be in error. (The depth is approximated by calculating the depth of the drilling bit when the cuttings started their journey to the surface.) A geologist may therefore be presented with cuttings, and an indication of the depth from which they came, that bear little resemblance to the properties of the undisturbed rock.

The geologist examines different samples of the cuttings and provides the experts with the associated descriptions. some of the samples possess all of the features expected for Berea sandstone, some have none, and some are ambiguous. Under these circumstances, the expert cannot be sure of the answer. The experts best opinion, however, requires knowledge of how the drilling process can cause the cuttings (as described by the geologist) to exclude some expected feature and include some unexpected ones.

The are many different effects on the geophysical logs, caused by the environment in which the logs are run. A count of the gamma rays, detected by the density tool at a knowledge depth in the well, gives a reading that is mainly dependent on the density of the mud in the borehole, the rugosity or roughness of the borehole wall, and the strength and associated statistical variations of the emission from the gamma-ray source. The log analyst is familiar with many of the techniques, and their limitations, used to “correct” away these effect when attempted to obtain a “true” value for the formation density.

Again, based solely on the logs and information provided by the log analyst, the expert could not deduce the depth interval of the Berea sandstone. Corroborative evidence is needed. The expert must know enough to tell the log analyst if the “corrected” readings of the density log are consistent with the density of the rock, calculated from the geologist’s analysis.

To convert a geological analysis into expected values for the density of a formation, the expert has to be familiar with the use of some basic physical equations. For a unity volume of rock containing solid and fluid material, these equations apply:

Fluid Density x Fluid Volume + Solid Density x Solid Volume = Rock Density

Fluid Volume + Solid Volume = 1

Starting with the “behavioral” approach, the expert can expect any Berea sandstone to have density values lying between the density of its lightest and of its heaviest constituents, respectively. These limits would be 1.0 gm / cm^2 (assumed) for the fluid, and 2.65 gm / cm^2 for the heaviest solid – quartz.

The expert has also learned that Berea sandstone normally contains less than 20%, by volume of fluid, and less than 5%, by volume of feldspar (density, 2.551 gm / cm^2). With these additional constraints on the above equations, it can be shown that the values of a density log across the sandstone should vary between limits of 2.315 and 2.65 gm / cm^2.

That “...it can be shown...” needs a comment. The log analyst with manipulating equations to describe the behavior of various petrophysical (physical properties of reservoir rock) parameters. The geologist is familiar with many of the features (e.g. the percentage-by-volume limits of the various rock constituents.) that are used to place rocks in some particular class. The problem is they are not familiar with the implications of each others’ descriptions. To write “it cannot be shown” above means the expert has to show the consequences of one user’s description to the other user.

After the density log has been adjusted for as many extraneous effects as possible, the log analyst can make a hypothesis about the corrected rock density at each depth logged. The expert can use these hypothetical values to exclude certain logged intervals in the well from the title “Berea sandstone.” Any corrected log reading of a formation density that is less than 2.315 gm / cm^2 or greater than 2.65 gm / cm^2 lies outside the “expected” density of the subject rock. (That exclusion may be wrong if the log analyst’s corrections were wrong. conversely, rocks that are not Berea sandstone might be included as “possibles” if they happened to have densities in the expected ranges.” The expert has only obtained a possible identification of the rock from the density measurements and the analyst. To improve the reliability of the identification, the samples of cuttings must be examined by the geologist.

The geologist identifies some samples in which the quarts grains, grain size or grain-contact type suggest Berea sandstone. Because of mixing or filtering of fine-grained material, a few samples contain features that are not normally found in Berea sandstone. Other samples are ambiguous because there is a absence of the constituents expected for such a rock. All these ambiguous cases,

together with the estimated depths from which they came, are relayed to the expert. The expert is now equipped with two lots of independently obtain, partly contradictory and partly supportive information. The users must then exploit the expert's knowledge to agree on the more likely depth intervals of the rock in question.

Much of the expert's knowledge relates to the processes that affect the users' results. The original questions had directed the trio's attention to a number of depth intervals in the borehole that might intersect with Berea sandstone. Some of the associations, however, are incomplete and possibly contradictory. To help resolve this conflict, the expert treats one user's description of the cuttings, or the log, as a correct hypothesis about the associate formation. Inserting this hypothesis into the associated processes, the expert can tell the other user what should be expected. The "expectation" appears in the other user's language: the analyst is advised of expected values for the density log; the geologist is told expected characteristics of the rocks in the cuttings.

Each user experiences a simulation based on expert knowledge. A good simulation allows the expert to translate one user's hypothesis into results, expressed in another user's language. Such a translation exposes one specialist's opinion, via an expert's knowledge-based skill in simulation, to the scrutiny of another specialist.

This point is fundamental to the substance of the paper and is worth restating. The expert's skill in translation, or as it is alternatively called, effect propagation, provides a most important facility to the users. Without being obliged to learn each other's specialized language and knowledge, one user can communicate with another through the medium of the expert. In this example, the expert is basically demonstrating the consequences of the geologist's hypothesis in the log analyst's language and vice versa. The expert does not impose a decision but, by refining the simulation as the users modify their opinions, provides the chance for the users to come to agreement.

Assume, in the example, the log analyst told the expert that the measured density of the rock was 5% too high, due to the effects of a heavy mud in the borehole. Remember that the expert knows the limits of density values for formations composed of Berea sandstone. As a result of the log analyst's revised list, the geologist realizes one of the "possibles" has obtained additional, supporting evidence. Unfortunately, the cuttings from this depth do not contain any of the expected fine-grained sandstone. The geologist asks the expert if that

missing feature can be reconciled with Berea sandstone and receives a snappy lecture on the efficiency of the sieving process, through which cuttings were recovered.

Where both hypotheses are supportive, there is associative evidence for the whereabouts of the sandstone. Where there is a contradiction that the expert cannot explain, the geologist can take another look at the cuttings to see if there is an alternative hypothesis. The expert can turn this into expected density readings and see if these readings meet the approval of the analyst. Contradictions between the geologist's and the log analyst's opinions are gradually removed as the expert's increasingly refined simulation is worked backward and forwards between them. The activity continues until results are obtained, satisfactory to all parties. Then, and only then, should the group decide on the most probable depth intervals for the Berea sandstone.

With time, of course, the expert learns many of the problems that affect results in the users' domains. The expert's simulations become increasingly sophisticated as more real-world processes are better understood. The sophistication is also essential if the expert is to answer users' requests for explanations.

2.3 EXPERT GEOLOGICAL INTERPRETATION

Questions such as "What?" "Where?" or "When?" can be turned into, "What thing, place or time?" The appropriate answer is a description of an object, or a feature of an object, that satisfies the question. Such questions are relatively unambiguous and readily understood by the expert. Everyday speech, however, is extremely loose in its use of words, especially in those questions which require answers in the form of explanations. Often, it is only possible to interpret the intended meaning of a word by the context in which it exists or, by the spoken emphasis given to the questions in which the word occurs.

The question, "What is fine-grained sandstone doing in these cuttings?" is an example of this complication. The right answer satisfies a questions like, "Why is there fine-grained sandstone in these cuttings?" Even then the meaning is not clear as, without an indication of emphasis, it is not possible to decide whether the questioner wants the expert to explain "Why...

...the sandstone in these cuttings is *fine-grained,"* or,

...the fine-grained material in these cuttings is *sandstone,"* or,

...the material in these cuttings is *fine-grained sandstone."*

The question, “Why...?” may imply the asker wishes to know the intentions of the party responsible for a certain state of concern. Alternatively, it may mean, “How...?” being a request to describe the process which resulted in that state of concern.

Some years ago, I witnessed the difficulties a young geologist had in giving a plausible geological explanation for opal fragments in a sample of rock cuttings extracted from a depth of 7,000 feet down an oil well. This caused great amusement to the roughneck who had seasoned the sample on the way over to his victim’s laboratory! Eventually, the geologist tumbled to the answer by giving up on the initial, “Why” meaning “How did opal get to be in such an unexpected place?” (The more insightful question, “Why is this roughneck laughing himself silly?” proved to be more productive.)

Back to a questions of the type, “Why is there fine-grained sandstone in these cuttings?” and let us exclude comedy as a possible explanation. In the absence of spoken emphasis, assume the questions means, “How did *fine-grained* sandstone get into these cuttings?” The expert’s answer depends very much on whether the “fine-grainedness” is expected or unexpected. If the cuttings came from a depth where fine-grained sandstone was unexpected, the answer might be a description of the process (e.g. mixing of cuttings from a nearby zone) that had the fine-grained feature as an associated, subsequent result. conversely, if rocks with these features were expected at that depth, the appropriate explanation might be a description of the depositional process (e.g. low-energy, fluvial) which produced the feature in questions.

To unravel what a user means by a questions, requires the expert to do more than just recognize emphasis or voice inflection. In fact, the users’ needs and expectations have to be included in the expert’s simulation of reality. Users do not spend much time qualifying their requests in order to remove ambiguities and make the expert’s work easier. They expect the expert to understand what they want to know and, perhaps equally important, what they do not want to know. The object or process selected for description, and the degree of explanation required, seems to depend on the expert’s perception of the user’s needs. One user may bet the following explanation:

> “The cuttings are from a Berea formation and therefore are *fine-grained* sandstones.”

Alternatively another user, asking exactly the same questions, may get a different reply:

> "Berea formations are laid down in low-energy, fluvial environments and therefore are *fine-grained* sandstones."

The above replies present the users with associations of features with behavior patterns to describe objects and processes. The first statement indicates that because the formation has the features of a Berea location, then cuttings form it inherit its fin-grainedness feature. The second statement describes a particular consequence of a more general process. It indicates that the time-dependent process of sedimentation (in this case, a river laying down sand particles onto what eventually becomes a Berea formation) results in grains that are fine in size, whenever the river carrying the sediment has a low energy. Which of the two replies the expert selects is determined by the perceived needs of the different users.

More of this in Section 3, but what is of importance to understand now, is the way in which human expertise can be, and is, used in practice. Obviously the expert knows a lot of the features and behavior patterns associated with particular objects in the domain of interest. This knowledge is also organized in such a way that it can be easily accessed by the user; the expert can infer what the user needs to know about the world in question.

Effect can be propagated, and explanations satisfactory to the user, given. The expert can tell the user "If you tell me these cuttings are from a Berea sandstone, then you can expect the grain size will be small." If asked "Why?" the expert can give a description of the process that is associated with the feature of "fine-grainedness." Of a number of possible descriptions, the expert can select the one most appropriate, in style and substance, to the particular user asking the question.

The foregoing examples illustrate the image presented by an expert to the users of geological expertise. The examples suggest some of the more important properties a system would have to possess, if it were to present a similar image. An expert seems to display five major qualities when interacting with others; learning, understanding, the ability to propagate effects, to handle conflicts and, lastly, to justify or explain results. These qualities have been translated into rather general specifications for expert systems and are presented below with appropriate comments.

2.4 SPECIFICATIONS FOR EXPERT SYSTEMS

If any expert system is going to satisfy a significant proportion of a user's expectations from a human expert, the designers and builders of that system must attempt to meet most of the following specifications:

1. Learning. Knowledge is to be inserted into the system directly, without using programming in the conventional sense. Human experts learn from their users, and can be taught in the particular language of each user. In this way, humans may accumulate a large number of facts or associations, but their expertise is distinguished from a mere catalogue of knowledge by the functional way that knowledge is organized. It has to be possible, for example, to have a expert system use an equation or a relationship like a human. (Such equations are not solved in the mind, at least not if they are at all complex. What happens is the expert knows the equation, as a form of shorthand, using it to describe the approximate behavior of interrelated variables.) The system has to learn how facts and functions are used, which is a different problem from their mere storage and retrieval.

2. Understanding. Interaction with the user must be in the language of the user. In the case of petroleum geology, a geologist would communicate with the system using geological terminology, and a log analyst, with the mixture of algebraic and numerical symbols that characterize this discipline. The system has to understand the significance of each interaction and reply in a language that is understood by the user. Human experts do not force their clients to learn or express themselves in a language that is convenient to the expert, experts accommodate themselves to the users. Experts are able to understand the user's meaning even if the messages sent by the user are not explicit. To decipher implicit information, the expert must understand the concepts with which the user is working.

3. Propagation of Effects. The consequences of one user's hypotheses are propagated into the domains of interest to the other users. An expert in petroleum geology is able to answer a questions like, "What would be your revised interpretation of this rock if you know the original density measurements were 5% too high?" Conversely, the expert could predict the value of the density measurement, if a geologist advised that a closer examination of the cuttings suggested the formation was composed of pure limestone. In the example given earlier, an expert was acting as an intermediary between a field geologist, examining some cuttings, and a log analyst trying to

correct a density log in order to give the best values for the rock's density. This ability to act as a kind of knowledge broker between specialists, is a key aspect of the human expert's repertoire of skills. The systems must also be able to play this role, of presenting to the users a clear picture of the consequences of their own, or each others' suggestions.

4. Conflict Handling. Frequently, a conflict occurs between real-world effects and the effect propagated from user hypotheses through the expert's simulation. The system must be able to recognize that such conflicts exist and give advice or initiate plausible, conflict-handling techniques. This activity, in the real world, involves close communication between user and expert: sometimes a user recognizes a conflict, decides how big a problem is represents, and act accordingly; sometimes it is the expert who recommends a particular solution, and the user has to decide whether to accept or reject. In either case, the parties have to be aware of what the other is doing and why.

Conflict handling involves complex pattern-recognition abilities and value judgments form the user. The specification for what an artificial system should do in this area is intentionally rather loose. A conflict, resolved by a system alone, would be a decision. Experts should not (and in my experience the best one do not) import solutions on their users. They examine alternative approaches to resolving conflicts and make recommendations, leaving the user to decide on the "best" course of action.

5. Justification. Results should be explained on request, and guidance given by the system if it does not know enough to answer the users' questions. Human experts can do this; so must any system that is designed to give users and equally convenient and accountable environment.

These five listed attributes can be described as either, the qualifications a recruiter of experts would require a human expert or, the specifications which an expert-system designer would have to satisfy to give a competitive service. It is now time to develop a theory of expert thinking or, how human experts are able to do what they do. Following the same sequence as the list of specifications, we start with a study of expert learning, or the acquisitions and, more importantly, the organization of domain knowledge.

SECTION 3: A THEORY OF EXPERT THINKING

3.1 LEARNING: THE BEHAVIORAL AND CLASSIFICATORY METHODS

Peter Sellers, in a scene from one of his Inspector Clouseau films, entered a hotel lobby in pursuit of his quarry. Faced with a rather large dog, lying near the stairway, he asked the proprietor whether his dog bit.

"No," came the laconic reply.

On patting the apparently inoffensive beast, Clouseau felt its teeth sink enthusiastically into his hand. Spinning round, he remonstrated, "I though you told me your dog didn't bite!"

"That's not my dog."

This cautionary tale illustrates the importance of placing objects in the right class, before applying expected rules of behavior to them. Interpretation in petroleum geology is faced with a similar challenge. At the risk of an over-generalization, it seems that the contributory sciences can be divided into two major groups according to the way in which the knowledge, associated with either group, is organized and used. This division of method is commonly observed when specialists, from either group, analyze geological problems. That does not mean it is a constraint on how they do their analysis but it does bring out a distinction between an expert's use of knowledge and that of a specialist.

Before taking this expert-specialist distinction further, it is worth considering how specialized knowledge in petroleum geology, as in most of the natural sciences, is divided into two relatively exclusive, schools of thought. I will call them the "Classificatory," and the "Behavioral" schools, respectively.

<u>Classificatory.</u> The classificatory school would include geology, mineralogy, paleontology, and sedimentology. The organization of this school's knowledge reflects the way in which that knowledge is applied to problems of interpretation. Emphasis is placed on a analytic procedure that goes broadly along the following lines. A precise description is attempted of the rock, mineral, fossil or sedimentary structure under consideration. The feature of this described object are compared with the characteristics of well-known objects, called typical examples or prototypes. These prototypes would have been

organized with a taxonomy, and allocated class names, associated with the features, or relationships, they have in common with other members of the class.

The classes, containing the prototypes, most similar to the object being studies, are examined to located a best fit according to hopefully well-established criteria of what constitutes “best.” Typically, some but not all of the features of the object would be available for measurement or evaluation. An inference is then made that the remaining, unevaluated features of the object studied, are similar to the equivalent features of the prototypes in the class “best-fitting” that object.

If the fit is not “good enough,” again based on some hopefully well-established criteria, an attempt may be made to find out how the object obtained the features excluding it form a class. The scientist may decide if the “anomaly” is sufficiently important to justify the creation of a new class. Alternatively, the anomalous features may be added to the list of features of the class that provides the best fit, so stretching the class to accommodate its more indigestible neophytes.

<u>Behavioral.</u> This school would include log analysis, statistics, seismology and geophysics. In devotees show a marked difference to the classificatory school in their approach to the problems of interpretation in the petroleum geology. From a theoretical or empirical point of view, an equation is developed relating several variables of petrophysical significance. Sets of such equations are used to solve for unknowns or, in the case of multiple solutions, probabilistic methods are used to derive mathematically optimal solutions. (These are the techniques commonly used in current, computer-processed interpretations.) The behavioral school (so named because what is considered is the way some properties of an object behave, as other vary) generally presents its results in the form of calculated values of critical properties such as porosity, permeability and hydrocarbon saturation.

A conflict may occur between the value of a property, as predicted by solving the equations, and a value of the same property derived from independent measurements. Two alternatives are then available to the behavioral school. Either the interpreter can change some of the input parameters to the equations or the structure of the equations can be changed to removed the conflict. (In fact, there is a third alternative which is to call the difference between what is believed to be true and what the process predicts, an irreducible error, and present the results to a certain degree of accuracy.)

The weakness of the two specialized approaches in acquiring and organizing appropriate knowledge should be measured in terms of the usefulness of the resulting interpretations. The main use of an interpretation is to find an accurate and economical way of describing geological objects and the processes in which they may participate. To construct these descriptions requires access to a suitable ordered knowledge based containing a sufficient range of facts and associations. That is sufficient is determined over time, by experience, and whether or not a new word or equation meets the criteria for general acceptance set up by the users.

Using a purely classificatory nomenclature, a geologist finds it fairly easy to talk about the grains of a sedimentary rock as “very-fine size, oblate shape, frosty textures, sutured contact type and so on.” Such a description can only be communicated effectively to someone who has a shared experience of the features of objects so described. (Otherwise, one man’s suture grain contact could become another man’s poison!) Assuming whoever is listening to the description can make these necessary association, the “essential” information about the features of an object can be transferred verbally. Problems start to occur when the features of an object vary significantly in space or time. For example, if a geologist used only feature names, it would be hard to describe accurately, how the flow rate of oil from a complex reservoir rock varies with depth.

In principle, an equation relating the behavior of the flow rate to a function of time, pressure, fluid viscosity, etc., should describe such a process more efficiently. Nature does not prove particularly helpful here. The field of petroleum geology is littered with the carcasses of formulae that started life as universal rules, principles or laws but expired in the hot sun of a real, complex world. The behaviorists have an excellent language to describe how the features of objects should vary with respect to each other. There is, however, a tendency to establish such a relationship for a particular object and then apply the same relationship outside the framework of proper theoretical and evidential support. The result is a body of knowledge, or to put it another way, of behavioral expectations that frequently fails to describe a context-dependent reality.

Experts appear to bridge the two extremes by tagging classes, objects in classes, and features of objects with what might be described as “atomic” behavior patterns. They avoid the trap of assuming empirical equations apply throughout a universal domain and, also, of trying to describe continuous feature variations with discrete feature names. In essence, the organization of expert knowledge

reflects the distinctive way in which that knowledge is used; to allow the assembly and simulation of complex, hypothetical, geological objects from more basic elements.

When an expert learns something new, the something is added to an established and related body of knowledge, efficiently organized to receive it. At this stage of the paper, it is not appropriate to say exactly what that efficient organization might be since it is a function of how the expert operates as a whole. For the time being, it is claimed experts absorb and store knowledge in the form of associated elements. One type of element is the features that have been recognized as useful in classifying, and distinguishing, the objects to which they belong. the second element is the associated descriptions of the way in which the features of these objects vary with respect to each other or, with the features of objects in other classes. Each description of a object (seen as a set of features) is then tagged with a description of a process (seen as a bundle of behavior patterns) so permitting the expert to link one object with another.

Expert learning may therefore be summarized in the distinctive way absorbed knowledge is organized. Specialists seem to fall into two, relatively-exclusive schools in terms of their knowledge acquisitions and organization. Neither school, alone, proves adequate to describe real and complex, geological objects. Each school has problems communicating with the other because their knowledge of real-world objects is structured so differently.

Experts avoid some of these difficulties, by breaking up specialist knowledge into elements that can be associated in a more flexible and natural way. Flexible associations of elements of knowledge allow better representations of real-world, geological objects and, also, give the expert the ability to communicate between specialists. Such representations leads us on to a discussions of the second expert qualification. How might knowledge be located in memory, activated, combined and manipulated in order that a state of “understanding” is created in the expert’s mind?

3.2 UNDERSTANDING: THE THEORY OF THE FUNCTIONAL IMAGE

The aim of an interpretation is to provide a functional description (discourse intended to give a mental image of something experienced) to someone. This someone requires the description in a form that can be used as a result or as a component in a further interpretation. The objects (the things which form elements or constitute the subject or a investigation or science) of study include

rocks which may or may not contain fluids. These objects possess properties *(qualities or traits belonging and especially peculiar to individuals or things).* These properties, and here is an important second definition from the dictionary, are *effect that the objects have on other objects or the senses.*

The last two definitions of a *property* bring out an interesting point. What the dictionary is saying is that, in common usage, the same word is used to mean either a quality belonging to an object or the effect that quality has on another object. the implication is not that we equate an object's qualities with its effects but that we have no way of distinguishing between them. It is not possible for humans to experience the properties of an object directly; the properties can only be interred from their effects on the mind. The mind may, of course, possess properties, but it is only aware of the mental objects in its possessions by virtue of their effects on itself.

I am therefore going to propose the law below, define some of its terms, and try and demonstrate form examples how human understanding may be considered as an instance of the law's application.

The effects of a system are isentropic with respect to its properties.

Entropy, in the above sense, is a measure of the information content of a message that is based on the logarithm of the number of possible equivalent messages. The name and visual image of Peter Sellers playing inspector Clouseau, represents a great deal of information. There would be few variations in the name or image that conveyed the equivalent information. These messages would therefore have a low entropy for the comedian's fans. Whether the dog that bit Clouseau was you or old, spaniel, dachshund or schnauzer is practically immaterial to the scene. There are a large number of possible configurations to the visual image of the dog that would convey the same amount of information to the film goer. The representation of the dog is therefore a high entropy message to the viewer.

The law is making a relativistic statement. A system has to have an environment (which may include an observing mind) before it can be "known," in the sense of experienced. Humans, inevitable, make subjective decisions on what constitutes a system or an object. They pick out a few of the near-infinite number of messages being broadcast from something e.g. the "doggy" features of the object that bit Clouseau. AT some vaguely defined point in time, and with even vaguer criteria, they associate a set of messages together and start treating the

thing sending the messages as a system. The thing is not aware it is a system: the mind affected by its messages treats it as one.

In this paper, and in the law of system isentropism, systems and objects are considered as things known only by their effects on the mind. Anything that has an effect on the mind may be deemed an object by the mind, that is to say, it is understood as an object that is distinguished from its environment. (An important point: "knowing" something, in the sense of mentally experiencing it, is different from "understanding" something, in the sense of extracting a meaning form it. When applying the law of system isentropy to the mind, the word "effects" refer to the first stage of what is experienced – before the mind makes an interpretation of what is meant. ***In the study of cognitive psychology, this is called the "sensory register."***

To obtain the meaning of an object, like a word or a sentence, it is not necessary to go down an indefinitely-long tree of definitions, more would such a procedure work. Apart from not being necessary, humans obviously don't do it. When we are familiar with a word, we do not stop to define it as it occurs; we somehow "use" it in the sense for which it was intended. Similarly, when we drive along a familiar road, we experience many non-verbal messages that have visual, tactile and audible dimensions. Only a few of these messages are assembled into objects that are understood as requiring action on the part of the driver.

This suggests that there are two stages through which information from an object has to pass before it can be incorporated into the nature of things which are understood by the human mind. The first stage is the use of the incoming information to located and activate remembered elements that are able to participate in the second stage. The seconds stage involves the use of these elements to construct a mental object having the same mental effects as the real-word object of study. AT this point ,the law of system isentropy can be applied. *If the mental effects of the two objects are the same, then the properties of the mental and the real-world object are isentropic. As far as the observing mind is concerned, the real-world object and the mental object have the same information content.

Now the properties of the mental object also belong to the mind containing that object. However, in an "understanding" state, the system (normally the human possessing the mind) can no longer distinguish between its experiences of the real-world object and the counterpart mental object it has constructed. This leads to a statement of the theory of the functional images as:

> *A system can be said the have "understood" an external object when it can construct an internal, functional image that produces the same internal effects on the system as does the external object.*

The external object can be any of such things as a tree, the scene of a dog biting a man, a command, a statement or a questions and to exercise the theory, it is worth going through a few examples. Let us take the liberty of imagining the world before its was every studies by the human mind. Then, taking an even greater liberty, make an anthropomorphic stab at the feelings of the inanimate objects, inhabiting this world.

Our world is a simple one with a meteorite falling in a gravitational field towards a planet. Already, this world is filled with three objects of potential study: a meteorite, a gravitational field and a planet. What are their feelings? As collision approaches, do they wince rockily? The answers is probably "no," but what is the meaning of our ability to understand the questions (which is a world object in its own right) and, equally important, imagine that the answer might be "yes."

An explanation appears to be that we have enough functional primitives in our minds to construct imaginary worlds that, whilst scientifically improbable, are entirely plausible to our senses. Such a world might be described by the predicate of the questions i.e., "The meteorite winces as it approaches collision with the planet." We can understand the question if we can, metaphorically, equip our imaginary world with properties so that its effects on our minds do not contradict this predicated description.

We know that humans are comfortable with metaphorical expressions (although there are restrictions on when and how they are used). World or phrases used in a literal sense to describe one particular type of object, can be applied to another type of object, if some important features of both types of held in common.

The descriptive word "wince" is a case in point. The world meaning would rarely be communicated to a child by a definition. Animals and humans can wince in anticipation of a blow, without knowing the word. A child learns that a certain instinctive grimace made to protect the face is given a name, "wince." That is just an association. If the reader were asked what the word, "wince" meant, the relevant facial muscles would most likely be located, activated and the effects described together with typical circumstances that might have provoked such a behavior. An observer of a rock being kicked by a man who had just stubbed his toe, can understand the kicker intends serious injury to the

"wincing" kickee. Neither of these emotionally-charged scenes of colliding heavenly bodies or, of rocks being mugged, requires world or definitions in order to be visualized and manipulated in the mind.

It appears then that humans can readily construct mental images of inanimate objects doing something metaphorically, like wincing. For this particular metaphor to be communicable, the primitive that generates the sensation of wincing must be associable with the primitive with effects that include the names, "planet," meteorite," or "gravitational field." Notice the association of the sensation with the last object, the gravitation field, feels unnatural. The action of wincing is normally witnessed only when a solid, visible (and living!) object is about to be hit by something. Meteorites and planets are solid and visible, albeit inanimate; gravitational fields are insubstantial, invisible and also inanimate. The acceptability of a metaphor, and whether it can be understood, seems to require a degree of likeness in the features of the object to which it is applied. ***(In Linguistics, this is referred to as "selectional properties.")*** A gravitational field has simple too few features in common with the class of objects that wince for it to be associated with the process of wincing.

The above analysis provides a possible handle for the meaning of world such as "true," "false," "nonsense," and "unknown." If the mind's mental-object generator is able to located and associate primitive that produce an object, having duplicate mental effects to those of the world object, then a description of the former is "true" of the latter. "False" would describe the situation where the mental-object generator produces effects which contract the world-object's effect. "Nonsense" would apply to the inability to make a "true" mental-object and "unknown" to the inability to make a "false" one.

The answer to a questions would then describe how to move from one mental object to another, in the mind of the responder. The first object would be the responder's image of the real-world questions, constructed by locating, activation and associating the appropriate primitives. A description of this mental object is then equally "true" of the real-world question, which has therefore been "understood." Suppose the responder know the answer to the question. The answer would have had to have been withdrawn from the responder's mental-object generator in order to make the question "true," as a description of both the responder's and the questioner's mental states. Describing the withdrawal process to the questioner, creates a second mental object in the responder's mind. This second object can best be described as the answered question. The effect of the second mental object is to make the

description of the first mental object (the original questions) “false” as a description of the state of mind of the responder. Whether it becomes false in the mind of the original questioner, depends on whether the answers was properly understood.

An answer to a question is then withdrawn (perhaps withheld is a better word) to make the question “true” and finally allowed, in the sense of no longer withheld, to make the question “false.” The answer is communicated by describing that process, which, when applied to the mental-object generator, makes the associated question “false.” We are accustomed to applying world like “true” and “false” only to answers or statements, not questions. There is only one basic difference, however, between a questions asked of someone and a statement made by someone. The questions requires an output, making the questions itself false as a description of the mental states of both the responder and the respondee. A statement presents an equal challenge to the responder’s understanding but does not require an explicit response.

The summarize, I see the act of understanding as the process of constructing isotropic, mental images of real-world object. Incoming messages from the real-world object are used by the mind to help locate and associate the appropriate primitives. When these primitives have been assembled into a mental object, with effects that are indistinguishable from those of the incoming messages, the mind has understood the real-world object.

The implications of the above reasoning for expert system are most important. If a body of associable primitives can be engineered from an elemental knowledge base, then at least two of the five specified expert skills are potentially available from the system: the ability to learn and the ability to understand. If the system is addressed by statements and questions and consistently succeeds in constructing replicas of these inputs, then the systems has learned and understood the inputs. The even more demanding abilities, to advise when a statement is false or to answer a question correctly, are the next expert skills for study. These activities require a needs investigation of the role of primitive in functional imaging.

3.2.1 The Use of Primitives in Functional Learning

Remember that, if a system input was a question, then it was argued that the answer to the questions was a description of the process moving the functional image of the question from a “true” to a “false” state. If the system could make

the questions true but not false, then the system "did not know" the answer, but did understand the question. If the system could not even make the question true, as a description of its own state, then the question was "nonsensical" as far as the system is concerned.

Earlier in the paper, the word "element" was introduced and applied to the smallest unit of knowledge used by experts to construct complex, geological objects. Seen in this light, an element of knowledge could exist anywhere; in a reference book, carved on a tree or located somewhere in the expert's mind. The concept of a primitive, in this paper, is distinguished form an element, both by its features and its functions. A primitive is considered to be the smallest mental unit of understanding. When a primitive is correctly located and activated by a signal, it generates a replica of that signal. It may generate other, associated signals at the same time, but according to the theory of the functional image, the mind, containing the primitive, has "understood" the meaning of the locating signal.

A primitive may therefore have a large knowledge content, with many elements of knowledge. The same utterance (spoken or written expression) may function as a primitive in a number of different configurations. The word "meteorite" in the mind of a young student is a case in point. "Meteorite" may be understood as an object associated with a definition like, "one of the small particles of matter in the solar system that reaches a planet without being completely vaporized." For the student to understand word that is unfamiliar, its definition has to be assembled from an association of primitives (i.e. the smallest mental units with which the student is familiar and does understand).

The meaning of the same word varies according to the context in which it is presented. If the student were asked, "Can you spell meteorite?" the utterance "meteorite" would be understood as a primitive object-to-be-spelled. If asked, "What does meteorite mean?" the student might reply with the memorized, dictionary definition. In this case the word is being understood as a primitive object-to-be-defined. If the student is familiar with what meteorites are, and do, the reply could be the effect of treating the utterance as an object-to-be-described, and listing the object's features and typical extra-celestial activities.

Whether the utterance "meteorite" locates in the student's mind a primitive object-to-be-spelled, -defined, or -described, it has at least one effect in common (that is, if it has been understood). The mental object it helps to locate, activate, and assemble generates duplicate mental effects to those of the utterance. The

other effects are the associated signals which generate the replies. Of interest is that the replies to the last two questions might be indistinguishable to the questioner. A student may use exactly the same words to define or to describe the object in question. A mind makes a quantum leap in its ability to understand the real world when it graduates form using messages for the location and output of definitions, to using messages that located and describe a functional image of that world.

Much human thinking is non-verbal even if the results of that thinking are communicated verbally with detailed scientific explanations to suit the listener's needs. These results are frequently verbal descriptions of mental objects that have been built form non-verbal primitives. Mental objects, so obtained, resemble, in their features and behavior, their real-world and non-verbal analogues. this mental world allows the human minds to understand (by building a functional image of reality) an indefinite range of real-world situation without recourse to language. The addition of language, however, increases the number of primitives available to the mind and, correspondingly, its ability to represent and understand reality.

It is time to apply these ideas to some domain-dependent problems. We therefore study how the primitives of petroleum geology could be located (placed or found), combined and then manipulated to produce realistic geological images in the mind of an expert.

3.3 EFFECT PROPAGATION: THE ASSEMBLY OF CONCEPTUAL MODELS

Earlier, I discussed how problems of interpretation in petroleum geology, were viewed differently by the classificatory and the behavioral schools of thought. This difference is reflected in the way each school organizes its respective, knowledge base. The classificatory school possesses a vocabulary that has developed around feature-based descriptions of geological objects. Depending on its application, this vocabulary can become inefficient and ambiguous. It is difficult to apply discrete feature names, with any degree of precision, to the complex geometry of geological objects or the feature names, with any degree of precision, to the complex geometry of geological objects or the time-varying processes in which these objects participate. This ambiguity is reduced when the users have a shared experience of the object to which the term applies, but the vocabulary then becomes increasingly obscure to the uninitiated.

The behavioral school organizes its knowledge around the way in which measurable properties of the objects are expected to vary with respect to each other. The literature of geophysics, log analysis, seismology, etc., is heaving with equations that are, in principle, effective for precise descriptions of complex and continuous behavior. Unfortunately, realistic constraints on the applicability of these equations are required by nature by rarely used in practice.

It has already been argued that experts operate differently from specialists belongs to either school. Their knowledge is organized in an elemental way, by attaching behavior patterns to the associated features of classified objects. When these elements of knowledge, singly or in combination, generate duplicate mental images to the messages that located them, they are called primitives. It would be nice to examine how plausible, geological models are constructed in the expert's mind, from the associated primitives of geological class and behavior. As explained below, such an examination introduces some rather subtle problems for non-experts, mainly due to their lack of familiarity with the domain and its terminology.

The following criteria could be used for deciding what is a primitive:

a. A primitive is the smallest mental object to be understood.

b. An expression is not in a primitive state, if it cannot be understood without the location of primitives.

These criteria imply that readers, unfamiliar with the domain of petroleum geology, would not understand primitives in the sense used by an expert. Only frequent exposure to the terminology and its use on real-world examples, allow the expert allow the expert to convert associated elements of knowledge to primitives for functional imaging. Most of us do not have this exposure and so would have difficulty in following a direct analysis of an expert petroleum geologist's thought processes. Therefore, instead of looking at primitives directly, the imaging process will be examined in two steps. First, in step one, the expert's mental assembly of hypothetical, geological models is discussed. Subsequently, in step two, we will see how these models can be manipulated so that their effects are made equivalent to the real world's, so creating and "understood" condition in the expert's mind.

Separating the process into two step, may appear to beg the questions of how experts use knowledge to understand situations. I have argued, however, that the information content of an expert's functional image should be equal to that

of its real-world object but still has an understanding of some hypothetical object. The latter object might exists elsewhere (meaning it probably would, since the expert's knowledge is essentially pragmatic); it just doesn't happen to be the real object, currently under expert observation. That means, strictly from the viewpoint of information content, we may just as well study how the real world is put together, as study the expert's model of it. Proceeding, therefore, through steps on the analysis, we start with the building blocks of geological objects – the classificatory features.

3.3.1 Some Classificatory Features of Geological Objects

Earlier in the text, the word "feature" was applied to an object, or a system, and requires some amplification. "Feature" is used as a substitute for the more familiar term, "property," to emphasize that system do not experience each others' properties, only each others' effects. A "property" of something is defined as a "characteristic," in contrast to a "feature," which is defined as a "prominent characteristic." "Features" better convey the idea of one system addressing the properties of another system, but only experiencing their effects, if the "characteristics" of this other systems are sufficiently "prominent." The "dimensions" of a feature are then, the ways in which the object's corresponding property are measured. The results of these measurement are called "values," and they may be qualitative or quantitative.

The literature suggests that features commonly used to classify sedimentary rocks are divided into sever major groups, with some groups containing sub-features. All features, or their corresponding sub-features, are "prominent": they must have dimensions and associated values to impact upon the observer (even if it is not always easy to find a word to describe this impact). The seven groups, and typical examples of their dimensions and values, are listed below:

1. <u>Composition.</u> Dimensions are the relative abundance, by volume, of the different kinds of mineral grains, making up the rock. These dimensions are usually expressed in an adjectival form – the noun form is rare. (Our Berea sandstone might have been described as feldspathic; it would not be described as possessing a certain "feldsparness" or feldspathcity.") Measurements in these dimensions are normally values that indicate the fractional, or the percental volumes of the rock, occupied by a particular kind of grain material.
2. <u>Grain Size.</u> Measured in various, related dimensions, typically using the Udden-Wentworth or USDA scales. Ranges that lie between certain numerical values of the nominal (diameter of a sphere of equal volume)

grain size are given qualitative values such as cobble, gravel and clay, in order of decreasing size.

3. Grain Shape. With typical, qualitative values in brackets the common dimensions of the grain-shape feature are: smoothness (polished, greasy or frosted), roundness (very-angular through sub-angular, sub-rounded to well-rounded), sphericity (with quasi-numerical values), "type?" (oblate, equant, bladed and prolate).

   If the typical, qualitative, values for shape type could be arranged in some rational progression, then a single dimension would suffice to distinguish them. With complex shapes this does not seem possible. Alternatively, neologisms for the dimensions of shape type could be derived from a noun form of the values. Words such as oblateness, equidimensionality, bladedness, and prolateness would then express the degree to which a grain's shape conforms to a type.

4. Fabric. The International Tectonics Dictionary defines fabric as the "manner of mutual arrangement in space of the components of a rock body and of the boundaries between these components." This feature has sub-features of both grain orientation and packing. In turn, grain orientation has sub-features of both crystallographic and dimensional orientation.

   Grain packing has dimension of density and proximity with values measured, according to a technique devised by Kahn (1956). Another "fabric" sub-feature is the type of contact of the rock grains with each other. "Grain-contact type" has values, defined by Taylor (195), i.e. tangential, long, sutured, and concavo-convex. Like grain-shape type, there seems to be no familiar term for the dimensions of this sub-feature. Perhaps "intercalatedness" would be suitable to indicate the way in which the contact line, in a cross-section through two grains, reflects the intimacy of their contact. Otherwise, "tangentially," "length," "suturedness" and "concavo-convexity" would have to do.

5. Color. Dimensions are, typically – redness, blueness, and greenness. Examples of the values: very red, blue, slightly green.

6. Location. The "address" of the rock in question. This feature has many possible dimensions; geographical coordinates, with some values of longitude and latitude; depth, with values indicating the distance below some reference level; name of provenance, or the formation form which the rock originated, e.g. Berea. One complication – the name of an object's original location is often interpreted to mean the expected or desired features associated with that location. (Had all the Aylesbury duckling or Dover sole every consumed, mature in the locations claimed, the human population would have vanished up its own ecosystem. Happily, the towns' inhabitants need protection against neither overhead incontinence, nor shoals of starting flat-fish!)

7. Age. The length of time since a referenced feature has existed. Like the location feature, there are many alternative dimensions use to indicate age. A rock extracted from a formation by a process known as coring,

could be described as Jurassic, 150 million years old, dating form the age of cycads, or any other formula indicating the time elapsed since it was laid down in a sedimentary deposit. Note that, just like the location feature, the age of an object should really refer to a particular object. For example, a Jurassic core sample could be two weeks old. The object, as a sedimentary rock, has “fabric” features that are about 150 million years old. In contrast, the object, as a core sample, has “shape” features that only come into existence two weeks ago.

It is instructive to look at the features of the next level of concern to descriptive, petroleum geology i.e. the gross structures, within which the rock reside. As we are dealing with solid, heterogeneous objects, it is not surprising that similar features are use to describe rocks (continuing grains) and formations (containing beds). the following table shows a few of these related features and examples of their respective dimensions or their values (the latter given in brackets).

FEATURES	ROCKS	FORMATIONS
COMPOSITION	Shaliness, liminess	Shaliness, liminess
SIZE	(Cobbles, very fine)	(Massive, laminar)
SHAPE	(Polished, frosted, Sphericity (Oblate, equant)	(Smooth, rippled) Conformity (Lens, prism)
FABRIC	(Tangential, long)	(Gradational, abrupt)
LOCATION	Depth	Depth
AGE	(Jurassic, Triassic)	(Jurassic, Triassic)

The above terminology is not intended to be exhaustive. It does illustrate, however, that in spite of a difference in scale, solid geological objects are described in similar ways. Nature, impartially, provides mankind with an infinite number of objects to play with, each object having countless properties. Those properties geologists select for consideration, are the ones having features perceived as relevant and useful. A sedimentary bed is defined as “a distinctive individual layer in a sequence of sedimentary deposits,” and it is evident that the geologist decides that makes a bed distinctive. Choosing the feature values of very-fine grain size, or slightly-feldspathic rock, for example, determines which

objects are cut out of all the possible objects for consideration. The grains or beds that materialize as a result of that choice, are defined by the feature values that have impressed the geologist.

The features of geological objects are measured by qualifying or quantifying dimensions, or both. The classificatory school is particularly adept with the use of the qualifying dimensions: what might be called the indefinite and the discrete. The behavioral school tends to operate with the qualifying dimensions: what might be called the definite and the continuous. Before studying how nature (and, arguable, the expert) combines the two, we will look at the way behavioral patters of geological objects are view by a log analysts – a member of the behavioral school.

3.3.2. Some Behavioral Patterns of Geological Objects

.c3.3.3.2 Some behavior patterns of geological objects

The physical properties of rocks can be measured in many dimensions. The most common way of representing how the values of these dimensions vary with respect to each other, is to use an equation that relates the values symbolically.

As an example, Archie's equations relates S_w, the fractional volume of the rock fluid that is occupied by water, to R_f, and R_w, the electrical resistivities of the rock and the water respectively. Another term in this equation is $F_{\%}$, the formation factor. F varies with the porosity (the fractional volume of the rock occupied by fluid) and increases as the porosity decreases. F can be determined empirically, from the ratio of R_f, to R_w, when there are no hydrocarbons in the fluid. (With no hydrocarbons, all the rock fluid is considered to be water, to that $S_w = 1$.)

Archie's equation is as follows:

$$S_w^{\ 2} = (F \times R_w) /\ R_t$$

Such an equation can be "understood" in a purely algebraic sense, without any geological constraints. If the other, independent variables are held constant, increasing F or R_w, or decreasing R_f will increase S_w. The equation could also be rearranged so that F or R_w became the dependent variable. The solution for S_w could be positive or negative.

If the equation is understood in Archie's sense, the behavior of the variables has to be constrained by geophysical and geological knowledge, $S_{w,}$ F, R_w and R_f all have to be positive. The value of S_w lies somewhere between 0 and 1. F turns out to be dependent not only on porosity but on some complex and poorly understood function of the rock's pore geometry, eloquently called tortuosity! Generally, log analysts would consider R_f , the resistivity of the rock, as the dependent variable. If certain constituents, known as clay minerals, were present in the rock, the whole structure of Archie's equation would have to be changed to described the behavior of R_f as the "independent" variables changed.

These concepts are more easily understood if the text following, is viewed as an annoted dialogue between the reader and someone at the other end of a communications link. The link is successively connected to a mathematician, a log analyst and lastly, a geologist. The mathematicians has been told of an equation but is ignorant of any of its geological possibilities. The analyst is equipped with instruments to measure various rock and fluid properties of a range of geological samples. the geologists is equipped with a hand lens to study the samples after they have been measure by the analyst.

The mathematician advises a message will be sent declaring sets of values of some variables, $S_{w,}$ F, R_w and R_f (which satisfy the equation) and asks the reader to find out what is "meant" by the message. The reader may be able to learn all the communicated sets of values, perhaps by ordering the sets into classes. The mathematicians might then test this "understanding" by giving three of the values in any one set and asking for the fourth. If a correct reply is only made when the three values belong to a previously communicated set, then it can be inferred that the reader has understood only in a classificatory sense. If a correct reply is received for any three values, the reader's understanding is at a behavioral level (The reader's primitive in this last case is a mental object that behaves in such a way as to generate the correct fourth number when advised of the other three.)

When the log analyst takes over the conversation, a first questions might be, "What do you know about rock resistivities?" The reader's answer is, "very little." Assume, after a chatty exchange, the two parties establish agreement on the meaning of the symbols, and that the equation developed in the head of the listener is given its proper name, "Archie's equation." The same problems as those given by the mathematician would elicit the same answers, expect the log analyst would have to tell the reader, "negative values of S_w and the other variables are not permitted." The reader's modified mental object now behaves

in a similar way to its predecessor. It does, however, have features like a name, Archie's equation, and a constraint not to output negative values when communication with log analysts.

The analyst continues sending along the values but for some of the reader advises, "I don't understand this message." If the analyst is told which sets are understood and why, it is possible to start classifying the rocks into two groups. One group behaves according to Archie's equation; the other does not. When the geologist takes over the conversation and examines the samples, the reader learns the division into "Archie" and "non-Archie" rock samples is related to the absence or presence of clay minerals in the electrically conductive part of particular samples.

An association can be made that for clay-free rocks, Archie's equation may predict fourth values or, if a rock has values satisfying Archie's equation, then it is clay-free. Varying the structure the "master" equation, the experimenters then find they can obtain appropriate equations to match the values of each geologically classified, non-Archie sample. Further, some of these later classification include several samples, described by just one of the modified equation. Such an equation can then be used to predict a sample's residual value, if the sample has been correctly classified and all but one of the values sent to the reader. Alternatively, if the sample has values that satisfy an existing from of the master equation, there is circumstantial evidence the sample belongs to the class, associated with that form.

The previous scene described how a long analyst might have derived equations, appropriate to the behavior patterns of an object's physical properties. The features of these objects allowed them to be classified by a geologist, and the behavior of the class members could be described by the analyst's equations. In the meantime, the reader had gradually built up expertise in connecting the two sources of knowledge. How, with a new sample, if the geologist advises its class, or some of its features, the reader should be able to tell the analyst the appropriate equations. conversely, if the analyst provide enough of the samples property values to establish which equation is solved by them, the reader could identify the sample's class for the geologist. In reality, the information from geologist and analyst would be insufficient for a certain interpretation. The reader, no considered an expert, could still, however, help the users come to a plausible interpretation.

The expert has learned enough associated knowledge to assemble a number of hypothetical but plausible geographical models for consideration by the users. These models should be plausible, because they are based on the expert's experience, albeit second-hand, of many real-world sample. fortunately, there are powerful constraints on the real world that the expert can exploit to keep the assembly process manageable. These constraints are best defined making a dimensional analysis of geological models.

3.3.3 Geological Models and Dimensional Analysis

The real world, unconsciously, synthesizes compiles geological objects. Through these objects, it propagates effects that depend on the feature values of the object synthesized. Some of these effects are noticed and examined by human observers directly – or indirectly via their geophysical instruments. Other effects are no doubt propagated by never experienced by humans. The real, geological world does not, however, solve equations; it contains objects with features that behave in a way human sometimes describe quantitatively – with equations and numbers, sometimes qualitatively – with descriptions or gestures.

Equations have to be dimensionally balanced in order that they may be applied to real-world geological objects. So, also, do qualitative descriptions. It a process is to affect the feature of an object, for example by changing the values of a dimension of that feature, then it must include at least one behavior pattern, having the same dimension for the dependent feature. (In this paper, a "process" in understood as the name for a bundle of behavior patterns. This bundle describes how an object's "dependent" features change, when an "independent" feature of that object varies.)

Described processes seem to fall into three main categories, in terms of the particular object feature, upon which the object's other features are considered dependent:

a. Location,

b. Age, and

c. Feature, other than location or age.

<u>a. Location-dependent processes</u> describe how an object's features vary, according to their relative position in space. "Grain size decreases upwards" describes how one of the features of a rock, its grain size changes with respect to variations in another feature, the location of that grain size. This particular

behavior pattern is likely to occur in a rock if its sedimentary grains were originally deposited in an old, river channel. "In an old, river channel," would then be the location feature's value, as measure in a second dimension. The first dimension implied, was depth, with a value, "up(wards)"; the second dimension implied, is provenance, with a value, "old, river channel." A complete statement would then be, "in an older river channel, a rock's grain size decreases upward."

Dimensional analysis can be used to ensure the expert's models correspond to what is "possible in nature. The expert must determine which processes can be attached to which objects, during the model assembly process, and which connections are impossible. Consider the previous example. Assembles models must describe what real rocks "might" do in old, river channels and therefore, dimensional integrity must be maintained. In nature, a grain-size behavior pattern has to have certain dimensions; i.e. the dimensions of the change in values of the grain-size feature versus the change in values of whichever features is being treated as independent. (If the grains size of a rock decreases upwards, in a partially-metric channel, then the behavior of the grain-size feature must have to dimensions of "change in grain-size. per unit of distance," e.g. "largeness," or "millimeters" per "foot.")

The process "grain size decreases upwards" is, in this case, a single behavior pattern. It can be attached to any object that has a grain-size feature and a location feature, e.g. a sandstone in a Berea formation, or a rock in a sedimentary bed. Notice we cannot say, "a rock's grain-size's redness decreases upwards," because the process of "redness decreasing upwards" does not include a behavior pattern with a dependent-feature dimension of grain size. We can, however, say "rocks redden upwards," because the object, "rock," can have a feature, "colour," measure in a dimensions, "redness."

In the real world, an object can only participate in a process, when a feature of the object and a dependent feature of one of the process' behavior patterns, share a common dimension. An object cannot be affected by a process that does not contain such dimensions; nature does not do it – an expert need not try.

Another useful constraint for model assembly is indicated by the direction of change, implied by the process under consideration. Moving vertically upwards in the sediments of an old-river channel, is apparently associated with a decrease in grain size. Therefore, sandstones should have a smaller grain size, the higher they are in such a channel and vice-versa. Similarly, if an old-river channel were described as the surface swept out by a parabola opening out as it moved

downstream, then a downstream location should have a greater distance between the channel's banks than an upstream one. The expert's models must satisfy these kinds of direction-of-change constraints: they affect the permitted, relative values of features when connected by behavior patterns that indicate a direction to the change.

b. Age-dependent behavior patterns would describe processes within which an object's feature values are events in time. Consider the values "fine," or "coarse," as used to indicate a sandstone's grain size. These values could be measured at moments in the grain-size's behavior when the grains were participating in the processes of "fragmenting," "grinding," "abrading" or "cracking." Notice dimensional integrity is secured since all these processes contain, by implication, a dependent-feature dimension used to measure grain size, i.e. "largeness."

Again, there is direction-of-change constraint on the real world and, therefore, on the expert's models. If the change of the grain-size feature's value is positive with the feature's age, then the younger feature value will be lower than the older one. Similarly, a "fine" grain size would always appear later in time, and therefore be younger in age than a "coarse" grain size. "Abrading" is a description of such a process, in which a rock's grain size is decreasing with time. If the word is legitimately applied to the real world, or properly describes the expert's model of that world, the younger size of an abrading grain is smaller and the smaller size is younger.

c. Other-feature dependent, behavior patterns would describe processes in which the "independent" object feature, is something other than a feature's age or location. We have already seen how such features of an object vary with others in our discussions on the application of Archie's equation. As before, these behavior patterns must induce dimensions for the change of the dependent-feature values against the independent-feature values.

The direction-of-change constrain also operates on objects that satisfy Archie's equation. If the equation is rearranged to make R_f the "dependent" variable, then it can be seen that the rate of change of R_f with F is positive. It for out that F, the formation factor of a rock, increases as the grain size decreases. Therefore, R_f must also increase with decreasing grain size, provided the other values are kept constant. (The process is more complicated in reality, but I prefer to sacrifice a little truth for a simpler explanation.)

Without recourse to number, nature would reveal itself as follow. If $S_{w,}$ and R_w are held constant, the rock, at the tope of the sediments in an older-river channel, would have higher values of formation resistivities of two rock samples were measured in the same, vertical section through the channel, then the sample with the greater value of R_f would be located above the other.

By appreciating how an object's features "can" vary with respect to each other, it is possible to build plausible models of reality. These conceptual models can be used as diagnostic or predictive tools. Let us assemble a more complex, geological model and see if it can propagate an effect without being supplied the answer.

3.3.4 An Example of Model Assembly and Effect Propagation

By adding to the expert's knowledge elements, we will try and construct a qualitative, conceptual model of a river carrying solid, sedimentary particles. As the river water moves down gradient, image the river channel alternately widens and narrows and the water speeds up and slows down accordingly. Assuming laminar flow, a plausible model, in the expert's mind, should predict an increased probability of sediment deposition as the channel widens; a lower probability, when it narrows.

Two of the features of a given droplet of river water are that the droplet has a location, with respect to the river source, and a speed in the streaming direction. The age-dependent behavior patterns for the droplet are as follows:

FEATURE	VALUE	BEHAVIOR	(SIGN)	VALUE (LATER)
Location	Upstream	Flowing	(+ve)	Downstream
Speed	Faster	Decelerating	(–ve)	Slower
Speed	Slower	Accelerating	(+ve)	Faster

Taking just two features of any cross-section of the channel, it has a location with respect to the source of the river and an area of cross-section, through with river water is running. Moving down gradient at any moment in time, the channel's area of cross-section acts as follows:

FEATURE	VALUE	BEHAVIOR	(SIGN)	VALUE (LATER)
Location	Upstream	Moving	(+ve)	Downstream
Area	Smaller	Widening	(+ve)	Larger
Area	Larger	Narrowing	(–ve)	Smaller

At a given point in time and space, a droplet of water in the river has a speed that is determined by the properties of the river at the same time and place.

River Speed = Volumetric Flow Rate of River / Area of Channel

Assuming a constant volumetric flow fate for the river, the equation above indicates an association of a faster water speed with a smaller channel area and also, a slower water speed with a larger channel area. Now, a solid particle, in suspension in the river, can be treated like a droplet of water, as far as its location and forward speed are concerned. Falling out of suspension is associated with slower water speeds.

The sedimentation process for the particle can now be described by the permissible association of a number of behavior patterns. These behavior patterns relate to changing feature values of different objects "at the same place and time." The process can be summarized below:

OBJECT	FEATURE	VALUE	BEHAVIOR	(SIGN)	VALUE (LATER)
Water	Location	Upstream	Flowing	(+ve)	Downstream
Channel	Area	Smaller	Widening	(+ve)	Larger
Water	Speed	Faster	Decelerating	(–ve)	Slower
Particle	Location	Higher	Falling	(–ve)	Nearer bed

Qualitatively, the above assembled model could predict results, purely by association. A particle is falling nearer the river bed, with the river water flowing

downstream, with a widening channel, with a decelerating forward speed for the water and the particle, the distance from the river source.

This last example described the assembly of a qualitative, geological model which behaved like simplified version of its real-world counterpart. The expert, however, can conduct object features to feature behavior-patters and crease highly complex and more realistic representations. Irrespective of complexity, these connections must still satisfy the same laws as the natural objects they are supposed to represent. The laws so far identified are: dimension integrity must be maintained; and secondly, features must have relative values that are consistent with the direction of change that are implied by the connecting, behavior patterns.

The manipulation of these increasingly-complicated conceptual models is difficult to imagine, when communicated by words. It is helpful to look at the problems of real-word complexity, and the associated demands on the model manager, in a more visual way.

3.3.5 Complex Models and the Bead-Curtain Metaphor

We will consider an extended visual metaphor for the process of model building and effect propagation. Imagine the model as being assembled like a bead curtain, hanging from a vertical frame. Each bead represents a feature of one of the object's modeled. The bead's colour represents the dimensions of the feature and the size of the bead stands for the feature's value. A number of grooves have been cut into each bead so it can be clipped onto an equal number of curtain strings.

Each string of the curtain comes with a predominant colour, but running through the string is a thread of another color. The strings represent behavior patterns of object features. The predominant colour of a string represents the dimension of a particular, dependent feature; the thread's colour represents the dimensions of an independent feature. String and thread are woven so that two "arrow" patterns (one of the string colour and one of the thread colour) are created on the surface of the string. The arrows point in the direction of increasing feature values, measured in the dimensions of the associated colour.

As an example, we could visualize a string to represent "abrading." The process, named "abrading" occurs when rough rock surface work on each other. The process includes and age-dependent behavior patterns of a rock's grain-size

feature. Let us assign the colour red to represent the grain-size dimension and blue for the age dimension. "Abrading" would now be, partially represented by a string with read arrows pointing in the same direction along the string as the blue arrows of the thread. (The larger the grain feature, the older the feature.)

With the conventions above, there exist some simple rule for building bead curtains which satisfy the demands of nature:

> Rule 1. Bead may only be clipped onto strings of the same predominant colour.
>
> Rule 2. For adjacent bead on the same string, the arrows of the predominant colour must point towards the larger bead. A string with no arrows of a predominate colour indicates no change in bead size; the behavior patterns of the dependent feature is conservative with respect to its associated, independent feature.
>
> Rule 3. Age-dependent behavior patterns are strings with the green arrows pointing downwards, towards the dependent feature's older values. The age of a features value increases downwards. The beads at the tope of the curtain represents the youngest state of an object's features.

The first rule secures dimensional integrity for the model. the second rule ensures that a feature' relative values are correctly positioned along the string. The third rule helps the model builder to organize object chronologically, with respect to the processes in which they are participating.

Imagine we decide to build an age-dependent model of an Aylesbury duckling. The raw materials for the model are some pieces of coloured string and a loose collection of coloured beads. Originally, the object of interest belongs to a flock of birds, cavorting in the skies above the town of Aylesbury. Subsequently, it becomes an object of gourmet attention, peering disconsolately out from under an orange sauce. (Even if the duck had never been to Aylesbury, it can still be treated as a canard!)

From the collection of beads, three colours are selected to represent three important features of ducks' weight, number of attached features, and shape (as always, we are short on words for the dimensions of shape, maybe "duckiness"

or “shapeliness” would do.) Following the rules of bead curtain building, we first clip the beads (representing the feature of the duck, on the plate) onto three different strings of appropriate, predominant colour. At an adjacent level below, and using the same three strings, we clip another set of three beads. The second set matches the colours of the respective strings, but its beads are not necessarily of the same size as the ones above. The new level, so created, represents the features of the duck in its older state, i.e. on the wing. All the age-dependent, thread arrows in the surface of each string point downwards towards the older feature values and away from the younger. The relative bead sizes for duck-on-plate and duck-on-wind now appear as follows:

OBJECT	RELATIVE BEAD SIZES			BEAD LEVEL
	WEIGHT	NUMBER OF FEATHERS	SHAPE	
Duck-on-plate	Smaller	Smaller	Smaller	Higher
Duck-on-wing	Larger	Larger	Larger	Lower

We are free to pick as many beads and strings as necessary, to adequately represent the duck. It is evident the object’s name, “duck,” may be applied to a number of associated features and behavior patterns. According to the context, an object may be eligible for a certain name, for a combination of a number of reasons: the existence of just one feature, the value of a feature, an association of certain features, or the behavior patterns exhibited by those features.

Can we put a name to the bundle of strings connecting the beads that describe the two state of the duck? It must be an age-dependent process that is associated with the three, contemporary behavior patterns of losing weight, losing features and becomes less duck-like in shape. Two candidates spring to mind, “plucking,” or “molting” and would be possible descriptions of what the bird underwent as it aged. (Both processes would require additionals to the basic model if the more likely had to be chosen. Extra features and behavior patterns would need to be spliced into the bead curtain for a more complete effect propagation. A “molting” duck is associated with an image of a bedraggled, crestfallen but still “living” bird. Such an image does not accord with the object lying on a plate with its legs crossed. “Plucking,” however, can be comfortable inserted in an age-dependent sequence, where that values of a bird’s feathers reflect the

processes of “flying,” “catching,” “dispatching, “ “plucking,” “cooking,” and “serving.”)

As a small digression, it is interesting to consider how models may be managed by using terms in their metaphorical sense. The changes in features values involved in a process called plucking, include decreasing either, fewer attached features and loss of “shapeliness.” A confidence trickster causes similar changes in the state of his “plucked” victim. “Plucking” could therefore be used meaningfully to manipulate, or to partially describe, an age-dependent model of the trickster in action.

“Fleecing” or “skinning” are synonymous with “plucking,” if used in this slang, metaphorical sense. Notice, however, that “molting,” a change in feature values like “increasing quantities of detached hair or feathers” would be associated with the literal object’s habitat. contrary to the other three process, “molting” is not therefore associated with increases in a crooked human being’s net worth!

Conclusion: metaphorical expressions, used to manage (manipulate or describe) a conceptual model must satisfy two conditions. One, the literal and metaphorical objects to which the expressions apply, must share the important dimension of the metaphor. Two, the feature values of the object, measured in these dimensions, must change in the same direction. (It cannot be said that “plucking” makes a victim “richer.”) Not unexpectedly, these two conditions are closely related to the first two rules of bead-curtain building.

The bead-curtain metaphor is useful, at least to me, when trying to visualize the management of conceptual models of age-dependent process. It has the great merit that qualitative effects can be propagated since only colours and relative, not absolute size and positions of the beads on their strings are considered. The mind can be “seen” creating a story or propagating its effects. A bundle of coloured strings, representing a process as a set of contemporary behavior patterns, is spliced into the curtain across a vertical interval. A set of beads of the appropriate colours and relative size representing an object as a number of simultaneous features, is clipped horizontally onto the string. Alternations made to the size of the beads, or to the rates and directions of change described by the strings, are transmitted through the rest of the curtain as they system obeys its rules. A horizontal section through the curtain gives a picture of the state of the model, at a give time, and reflects the effects of the causal chains, hanging below that section.

The bead-curtain metaphor continues to be useful when studying further complications to the model-building process. One such complication is that humans are accustomed to think of an object keeping its name, even though additions, subtractions and modifications are being made to its features.

In the single space of age-dependence, for example, the simultaneous feature values of an abject may be very different in age. A particular grain of sandstone might have a composition value older than its texture value. In turn, the grain's texture value might be older than the value of the grain's contact type with its neighbours.

Imagine such a grain as a set of beads, clipped onto the appropriate strings, starting out in life at the bottom of the bead curtain. The brain is born by erosion form a source rock and possesses a "composition" bead, indicating "made of sandstone." Other beads at the same level indicate relative values for the "surface texture," "contact type" and the "location" features of the grain.

As time goes on, the grain is "carried" down a river and deposited on a beach where it suffers the "abrading" action of the waves and its fellow grains. The bead representing the "made of sandstone" is conserved in size since abrasion does not change "composition." The grain's originally polished surface, however, becomes pitted and rough. One of the "surface texture" leads, with a colour that indicates the "smoothness" dimension, therefore, shrinks. (Conversely, the bead indicated the "roughness" dimension, expands.) The "abrading" process ends, when the grain is buried under some sedimentary deposit. The value of the "smoothness" dimension and the size of the associated indicator bead, is conserved form that moment on. Under the enormous pressures generated by deep "burial," the grain is forced into closer and closer contact with its neighbours. The "contact type" develops from "non-existent" through a "tangential" to a "long" value. The "location" feature in dimensions of "distance from source," varies from a "small" through "large" to a conserved value, as time increases.

The set of beads, representing the grain, can be considered sliding up the curtain strings, with increasing time and decreasing age of each feature value. New strings are added, beads change in size, some beads may disappear as the object participates in different processes. Throughout this time, the grain somehow kept in identity in the model-building, expert's mind. That is a major complication: which features must be kept together and which features maybe added, discarded or changed in value, without an object having to change its

name? The answer can only be found by experience and it probably varies according to the context.

Yet another complication of model building is that an object's features may vary with respect to other features besides age. If the side of a hill is cut away, the exposed face shows a two-dimensional picture of the sedimentary rocks. The picture is a result of depositional and burial processes that may have taken place over millions of years of geological time.

Imagine our long-suffering grain has just been exposed by the hill-cutting activities of some road builders. Time is frozen at the present and a feature set, called "the grain," is moved upwards in the exposed face, from its starting location in the bottom sediments of an old-river channel. The features in the set and the feature values, themselves, will change according to each new location. It will no longer be possible to refer to the object as "the grain." The correct name, for the modified set of features under examination, will depend on what and where these features are.

Traveling upwards in the old-river channel, the bead that indicates grain size is going to shrink, i.e. there is a negative direction to the change of grain size, with distance from some reference point below. The independent features in this upward-moving process is the "location" feature, with dimensions of "distance form some reference level below." Now, the size and the colour of the bead, representing a particular grain's size, are the same whether the bead is on a location-dependent string or, on a age-dependent string. The same features value of the same grain can be, and frequently is, considered as an event in several behavior patterns. The grain size might be described as a small because:

- the grain was a long time abrading (age-dependence),
- the grain is at the tope of an old-river bed (location-dependence), or
- the grains read a high resistivity (R_f -dependence).

Because a bead can be on an umber of strings, each beam must have enough grooves to cater for the strings that "can" be clipped into them. To represent different independent dimensions, strings of the same predominant colour need to be differentiated by the colour of their interwoven threads. Unfortunately, the relatively simple, bead-curtain metaphor is inadequate for modeling dependent features varying against multiple, independent features: it is too difficult to visualize simultaneous manipulations to the beads in more than two,

independent dimensions. Most of the original's imagery can be salvaged, however, as it is still relevant to the rules of conceptual-model building.

I have argued the expert stores knowledge in an "atomic" form. This basic form can been seen as a nuclear bead connected by a number of strings to other beads of the same colour. These other beads are larger than, smaller than or equal is size to the nuclear bead, if their connecting strings have outwards or inward points, or absent, arrows. The thread arrows continue to point in the direction of increase of the independent-feature value, and the thread colour indicates the dimensions of that feature.

How then, might an expert assemble a geological model with more than two independent-feature dimensions for the objects involved? Let us return to the problem of the river carrying sedimentary grains into a widening channel. Each object in the problem for example, river, channel and grain, pulls from the expert's memory sets of nuclear beads that represent the features commonly associated with those object. Each nuclear bead pulls with it a number of strings. Each string, as a behavior pattern, represent part of a process with which its nuclear bead is commonly associated e.g. flowing, decelerating, widening and grain size decreasing upwards. Trailing on the outer ends of the strings, are beads of no particular absolute size but, with a relative size determined by the string-arrow rule.

If individual grains are being examined, at least two processes can be considered to explain their relative size. The size bead has two string attached to it; one indicates beads get smaller as their location, in an old-river bed, moves upwards, the other string indicates beads get smaller as the age of deposition, at a given river speed, decreases. Through the same bead, therefore, there is an age-dependent string and a location-dependent string. If a more comprehensive model "allows" the presence of these two processes, the expert has access to two, plausible associated explanations, for the variations in grain size.

It is time to make a summary of the first step in the functional imaging process. I have attempted to described how an expert might go about assembling plausible models of complex, geological objects. The expert's models are plausible, if they satisfy the same constraints of dimensional integrity as the real-world objects they are meant to represent. The models are initially qualitative because they are built according to rules governing the relative, not the absolute, values of object features. An interpretation of an object is implied by a description of its modeled counterpart. Even if a particular interpretation is wrong, the model-

assembly procedure provides a coherent structure for a function, as opposed to a purely referential, knowledge base.

Much of human thinking appears to involve the manipulation of mental models. This is most evident when the thinker is imagining the possible consequences of a real, or hypothetical perturbation to some synthesized image of the outside world. The expert petroleum geologist must and does have the mental ability to assemble features into objects and, using the bundles of behavior patterns associated with such features, have these objects participate in geologically plausible processes. The activity requires a knowledge base so organized that object features can be inserted into a model with their associate behavior patterns attached and usable. Conversely, if the knowledge base is well-organized, dimensionally-correct and carrying comprehensive associations, it can be used to predict effects that are at least "true" of a realistic, geological model.

The problems of conflict handling are inextricable linked to the problem of understanding. Conflicts can occur when the expert has assembled and exercised a conceptual model of a real-world object. This model may produce, or propagate, internal effects on the expert's mind that contradict the effects of the external object being studied. In the expert's mind, there is simultaneously a state of conflict and a condition of object-not-understood. Step two, in the functional imaging process, considers a number of techniques for removing the conflict and reaching an "understood" condition.

3.4 CONFLICT HANDLING: THE PRINCIPLE OF RECIPROCITY

Returning to the metaphor of the bead curtain, the row of beads at the top of the curtain was described as representing the feature values of an object at their youngest state. This was an arbitrary constraint: a different choice of strings could have represented the state of an object at some point in physical space (location, as independent-feature) or in any other space. We could also insert time-varying, age-dependent strings, above the top row of beads, in order to represent future processes and possible events.

Let the curtain be made multi-dimensional, with an indefinite number of independent features. We can then extend it to represent many objects and processes before dropping it on the floor as some immensely complex mess of beads and string. Each bead and string must still obey the building rules, referred to earlier. The exception is that age-dependent arrows no longer have

to point downwards but they must continue to point in the direction of decreasing time (increasing age).

Suppose a hooked rod is used to bring out all the features of an object, at a given location and time, from the mess of the floor. The beads, representing the dimensions and the values of the object's features, are plugged into sockets on the underside of a glass table. Seen form above, the object is represented as a surface made up of many different sized and coloured beads. The expert can also see that each bead has strings of the same predominant colour, hanging from it into the mass on the ground. The strings indicate the way in which the features would behave if the expert slid along the age, location or other independent-feature thread. Bundles of the strings indicate the processes which brought the object to the state it is in, and took it to the state it later, or further, became.

Say the expert trod mentally on some of the beads on the floor and they were squashed, shrinking in size. Effectively, an experiment has been made with perhaps both the history, and the future, or the object modeled on the underside of the table. Provided every bead does its duty by its neighbor, and shrinks or grows as determined by the interconnecting strings, the effects of the experiment are propagated to modify the object at the table. The expert does not need to be aware of the detailed mechanics of the propagation but sees the picture at the table change in response to what has been done to the beads on the floor. The beads at the table shrink or grow, new one appear and existing one disappear as the model object responds to events in its past and its future. It is evident that is should be possible to learn the approximate response of the object for particular changes to the mass. Eventually, a "table picture" can be constructed on demand (which is what has happened in the reader's mind, if the above scene has been successfully communicated).

To successfully communicate something to someone implies the same message, intended for transmission, has been understood by the receiver. Understanding has already been defined, at least for the purposes of this paper, as "that property of a system which exists when the system has constructed a functional image, producing the same internal effects on itself as the object to be understood."

Let us continue with our furniture metaphor. The internal effects of a world object on an expert's mind can be visualized as the picture, generated by beads plugged into sockets on the outside of a glass ceiling panel. In this case, the

string run out through the sensory components of the system and into the real-world object of interest, whether it be a rock, a table or a question in a geologist's head.

There are now two pictures to look at, one in the panel above and one in the table below. If the expert wishes to understand the "panel" picture, it will be necessary to get busy with the hooked rod and place a few strategic kicks into the knowledge base and make a look-alike "table" picture. Which beads are initially pulled out by the hook, and then plugged into the table, would obviously be closest in size and colour to the one in the panel. Where the kicks are directed at the knowledge base, would depend on the expert's memory as to the expected effects on the table picture. If the table picture is successfully driven into agreement with the panel picture, then ,at least for the expert, a functional image of the real-world object has been assembled.

The real world is not in the habit of challenging the expert's understanding with a presentation a series of identical objects. Normally, real-world objects appear as slightly-perturbed versions of more familiar ones. Experts must therefore find the equivalent, conceptual perturbation to their more familiar, functional images. An "understood" condition, with respect to the perturbed objects, can then be recovered.

Given the state of "perfect" understanding, it is not possible for the expert, in the metaphor, to distinguish between the impressions of a real-world object and its mental image. Both have the same effects on the mind. Let a small and defined perturbation be made to the mass underneath the table. Provided that the model used to generate the mental object is continuous and deterministic, small but recordable differences will be developed between the panel and table pictures. Many such experimental perturbations could be made and recorded in memory together with the associated picture differences. The expert's memory could then contain three, associated lists:

1. The descriptions of mental objects when their real-world counterparts were well understood. (These descriptions are also interpretations of the real-world objects.)
2. Some defined perturbations to the expert's mental-object generator for each mental object in list (1).
3. The feature differences between the real-world and mental objects that result form each of the perturbations in list (2).

Imagine the real-world object is now slightly perturbed by nature from a state named OBJECT B, to a state named OBJECT A. The expert sees a subtle change in the original panel picture. “Understanding” OBJECT A, requires manipulation of the mental-object generator until identity is obtained between the ceiling panel and the table pictures. The first step in this process is to search the memory for a certain OBJECT B, and plug it into the underside of the glass table. OBJECT B is what is available in the expert’s memory that best resembles the perturbed real-world object. The features of the two pictures are then compared and their relative differences noted. These differences can be seen from the viewpoint of OBJECT B, looking up, from OBJECT A, looking down. In the latter case, the differences would be the reciprocals (inverse relationships) of the former. Examples are given below. The convention used is that the feature differences, listed in the OBJECT A column, are seen looking down from the panel; those listed in the OBJECT B column are seen looking up from the table.

RECIPROCALS	
OBJECT A	OBJECT B
Below	Above
Taller	Shorter
Faster	Slower
x times larger	x times smaller
More acceleration	Less acceleration
Later	Earlier
Closer	Further
Heavier	Lighter
Feature present	Feature absent
+ 0.02 difference	– 0.02 difference
Positive difference	Negative difference

The generator has also to search the memory for OBJECT B’s sets of feature differences, together with the descriptions of the model perturbations that cause them. (These sets were memorized after a real-world OBJECT B had become well-understood, and experiments were made to vary the state of its image.) Assume these memorized sets of features differences were recorded from the panel, looking down, i.e. using the real-world OBJECT B as reference.

Suppose a set of OBJECT B’s feature differences is found in the memory, identical to that recorded in the fight-hand column above. The view looking down, in the remember situation, is then the same as the view looking up in the new one. In

addition, the views are from the same object, OBJECT B. The perturbation, associated with the matching set of feature differences, is therefore a probable description of the perturbation made by nature. Making this perturbation to the model, should drive the table picture from OBJECT B towards OBJECT A. This process can be continued indefinites as sets, point to perturbed objects, point to sets, along a path of successive approximation. Complete understanding is not achieved unless a perfect match occurs. Convergence is likely if OBJECT A exists in a well-explored domain, and unlikely if the real-world object is unfamiliar.

To put this technique in a more obvious form a principle of reciprocity for “understanding” system can be stated as follows:

> *View from the same position inside an understanding system, a perturbation to a real-world object will generate a set of feature differences that are reciprocal to the differences, generated by the same perturbation to the object’s functional image.*

In practice, this principle of reciprocity could provide a power technique to recognize and understand new situations that are fairly closely related to previous experience. When real-world objects or situation have been well understood, purely mental experiments can be made to see what their perturbed images look like, viewed from the real-world. Nature may subsequently bring a perturbed object to the expert for examination. The object can be looked at from the viewpoint of the most similar, remembered and unperturbed image. If a set of feature differences is seen that is reciprocal to a set existing in memory, the associated image perturbation can be identified. Applying this perturbation to the undisturbed image, should reduce the object-versus-image conflict and move the system closer to an “understood” condition.

We have now reach a possible explanation on how experts handle users problems. According to this explanation, the expert assembles and manipulates mental models until a functional image of the subject world, including the user’s problem, is reached. If the answer is already known by the expert, the statement of the user’s problem is a description of only the user’s state of mind (the expert’s state of mind must be perturbed before the problem can be a “true” description of it). This perturbation has the same information content as the answer: it changes the statement of a problem as a description of the expert’s state-of-mind, from “false” to “true.” The expert’s description of the necessary perturbation appears as an answer to the user’s problem.

If the expert does not have foreknowledge of the answer, it will have to be derived. We know that the statement of the problem is made "true," as a description of the expert's state-of-mind, if the problem is "understood." We also know that it must be make "false" if the user is to get an answer. A considerable amount of expert model assembly, effect propagation and conflict handling may be necessary, using the techniques discussed, to provide this answer. The assembled, requisite model must satisfy all the constraints of the real world (dimensional integrity, for example). It must also have the effect of making a statement of the user's problem "false," as a description of the expert's state of mind. This effect, this time generated by a model the expert did not know in advance, appears to the user as an answer to the problem.

After this study of step two of the functional imaging process, my interpretation of the expert's use of primitives should be restated. A mental object, or conceptual model, is in a primitive state if it generates a duplicate of the message locating it. If the duplicate has to be derived by accessing other objects, then the derived object is not in a primitive state. The object may become a primitive, if the expert is frequently exposed to the same, real-world situation. The complexity of its information content may also belie the description "primitive, but the term is mean to convey what a mental object "does," father than what the object "is." A probably-indeterminate amount of expert knowledge exists in this primitive form; many problems are immediately understood through repeated appearances; many answers are pulled straight form memory. Whether mental objects are remembered or derived, an expert's services form a knowledge base depend on the quantity, the information content and the combined ability of the expert's primitives.

Of course, the user is interested in knowing why an expert came to a certain conclusion and will frequently ask for some justification. I believe a system's ability to explain itself is a natural consequence of its ability to construct a functional image of the real-world object that requires explanation. This idea is discussed further in the following pages. If it proves practicable, the last of the users' requirements from an expert system the justification of its own results, will have been realized.

3.5 JUSTIFICATION: THE EXPLANATION OF RESULTS

Being asked to justify an opinion about an object or process, is not very different form being asked to give that opinion in more detail. When the expert has an opinion about an object, a user can experience that opinion in two basic way – a

description of what the object was, is or will be or, a description of the past, current or future processes that affect the state of the object. According to what has gone before, the descriptions refer to the expert's model of an external reality. For the user to understand these descriptions, an image of the expert's model has to be constructed in the user to understand these descriptions, an image of the expert's model has to be constructed in the user's mind. If the user is unable to locate or construct an image corresponding to that model, then the expert's opinion requires justification. Expert justification is then seen as providing sufficiently-detailed information to complete the user's image of the expert's model.

Reverting to the furniture metaphor, the expert's model has been assembled from various features and behavior patterns to generate a "table" picture. The expert's opinion is meant to create a duplicate "table" picture in the mind of the user – let us assume this has been unsuccessful. The user's request for an explanation appears as a "panel" picture which the expert must duplicate to understand. The question the expert has to get the model to answer is, "what knowledge has to be withheld form the model for the expert, also, to need the same explanation?"

The question about what needs to be done to the model is not generally made explicit in the mind of the expert. Instead, the "panel" picture, representing the user's questioning state of mine, is used to drive the expert's model and generate a duplicate "table" picture. Remember, if the strings running through the beads under the table are pulled, independent-feature processes are activated and can be described. Similarly, as beads respond by growing or shrinking, appearing or disappearing, objects take up different states and can also be described. Manipulating both, the expert tries to recreate the user's problem by withholding the information which answers it. If successful, a description of what the expert withheld appears to the user as the explanation desired. Let us look again at one of the user's questions of Section 2, "why is there fine-grained sandstone in these cuttings?" Assume that a conceptual model has been assembled in the expert's mind which described the processes resulting in "fine-grainedness." Assume the user's inflection was on "fine-grained" so the expert can tell which value (fine) of which feature (grain size), of which object (sandstone) at which location (same as the cuttings) was in question.

A number of string through the "grain size" bead can be pulled in order to examine the processes within which the particular value "fine" is an event. If there are several strings, there are several processes. If they satisfy the real-

world constraints of dimensional integrity and direction of change, the processes are all plausible. Which ones are more probably, depends on their frequency of association with the objects affected. Withholding these processes from the model could provoke the expert to ask the same questions as the user; withholding knowledge allows the user's question to be understood.

Describing these "withheld" model processes could result in any of the following answers:

A. "Because the grain size in old-river channels decreases upwards, and these cuttings are from the top of an old-river channel."

B. "Because finer grains are deposited later in rivers, and these cuttings are from the late deposits in an old-river channel."

C. "Because the smaller the grain size, the denser the rock, and the density log shows the cuttings to have come from a dense rock."

(A), (B) and (C) include partial descriptions of location-, age- and density-dependent processes, each one being qualified by the circumstances in which it occurs. An association is being made between some of the features of an object and their behavior patterns.

It is therefore quite possible the expert delivers a different explanation to the same questions form a geologist, a sedimentologist or a log analyst, respectively. The explanations may look like non-sequiturs to the other users. The appropriate explanations would depend on the expert's assessment of what the questioner has failed to understand, i.e. the difference between the questioner's model and the expert's model. It is possible to have many alternative explanations; the "best" one would convey the required information to an individual user in the most efficient way.

To explain why a feature of an object has a certain value, the dimensions of the independent features, on which that value depends, just be located in the expert's model. To select the appropriate explanation, the expert has to assess which of these dimensions is missing form the user's model and provoked the original question. The expert's reply should provide enough information, neither more nor less, to make the question "false" as a description of the user's model. this information would normally be presented in two parts. Firstly, a description of the behavior pattern that relates the feature value in question to the missing, independent-feature dimension. Secondly, the circumstances in which this behavior pattern occurs.

With the expert's justification, an appropriate "nuclear bead" can be "spliced" into the user's defective model, in the correct position relative to the other beads. As a result, the user's model moves form a state described by the original questions towards a state described by the expert's model. Assuming a description of the expert's model is a good interpretation of reality, a satisfactory explanation has been communicated.

The reader may wish to extend the metaphor to justify processes from the states of participating objects. Hopefully, the principles of justifying an expert's answers is not necessarily a description of the model generating the state in questions; the difficult part is for the expert to model what the user does not understand about that state and express the explanation accordingly.

This paper started with a study of what an expert represents to the users of expertise. I then discussed a theory of expert thinking that attempts to explain how human experts are able to provide knowledge-based services to their users. It is now time to consider how a user might appear to an expert – of both the human and the "system" variety. This investigation looks at the extent to which an artificial system might provide human-level skills and, also, at some philosophical limits both the human's and the system's ability to "understand" their environment.

SECTION 4. HUMAN AND SYSTEM EXPERTISE

4.1 MULTIPLE MODELS OF SINGLE OBJECTS

Special techniques in knowledge organization provide some of the important distinctions between experts and users of expertise. Users were earlier categorized into either the behavioral or he classificatory schools for one particular domain: petroleum geology. An expert has to identify with which type of user he is dealing. This is more than just a matter of recognizing he different use of symbols and terms: the is a procedure difference.

The behavioral schools seems to operate form a catalogue of precise process descriptions, mainly in the form of equations. These equations are solved for feature values, treated as ends in themselves, rather than means to better classify the object studied. The classificatory school uses a library of fine distinctions between feature values to classify objects. Many of these distinctions are visual and tactile, but accurately measurable, and are linked together by only loosely-defined processes.

The above observations are, of course, simplifications. People are fee to address the problems of petroleum, geology in any way that is comfortable and effective. Frequently both techniques are exploited, but it is the expert who develops the skill of conceptual-model assembly to the highest degree. The expert's role is to assist one, or several, users to assemble mutually consistent models of the same geological object.

Short of achieving mutual consistency, the expert's model must propagate the effects of one user's opinions into another user's language. Put another way, the expert's model must be kept geologically consistent, whilst one user drives it around another territory.

These notions are subjective. A near-infinite amount of information about even a simple object would be required to convey a complete description of its properties to a human. In reality, only a few properties manifest themselves as features that can be experienced and measured by humans and their instruments. In turn, only some of these features are separable form other features and identifiable as relevant and useful to human needs. An absolutely "true" and complete interpretation of an object may be desirable, but it is probably impossible to prove that state was ever reached. The expert sees the users as filters for information form the outside world. The partial, subjective

information they give to the expert has to be worked up into a model that is described by the information given, plus the propagated effects. The users can then be shown that, either their opinions are consistent with the expert's model or there are contradictions they need to resolve.

An expert's world can be visualized with the metaphor discussed in the pervious section. A few users are feeding information and opinions about a real-world object into individual glass panels in the expert's mind. A mass (knowledge-base) of beads and strings is lying around in the expert's memory, waiting to be assembled into a table picture. The expert, using all the methods discussed, force the table pictures to be consistent with each other and tries to make them identical with their respective panels. If a contradiction occurs, the table picture is driven out through the offending panel above it. The user, operating that panel, is then presented with the expert's opinion based on expert knowledge and the other users' inputs to the expert's model.

The expert's opinion has come through a panel where the object's original features were through believed by the now-contradicted user. by the reciprocity principle, the expert's communication is automatically in the language of this contradicted user. Consequently, the user should be able to understand the contradiction and the circumstance where it applies. The user must now decide if the original opinion can be modified (to match the expert's) or whether it should be imposed on the expert's model. If imposed, the opinion goes back through the panel and focus the table picture underneath into identity. A "wave" of revisions sweep through the expert's mental-object generator to be reflected in changes to the other table pictures. These pictures are pushed out into the other users' world for consideration, accommodation or rejection. This kind of knowledge-based translation activity by the expert continues until users accept the consequences of each other's opinions about the geological object of interest. A description of the expert's model is then the interpretation most acceptable to all the users.

It is interesting to imagine an entropy map of these exchanges. Remember, entropy was defines as the information content of a message based on the logarithm of the number of possible equivalent messages. The entropy law states that a system's properties and effects (which can be considered as received messages) are isentropic.

A real-world object has a very low entropy since its "total" information context precludes any alternative configurations. Features of a real-world object that

actually impact or concern the human user will never contain this total information. Therefore, relative to reality, a conceptual model as a higher entropy value; different configurations of an object could send the user equivalent messages.

Any experts best interpretations is a description of the lowest-entropy model that is consistent with both the expert's knowledge base, and the opinions of the users. The best-expert's best opinion is whichever of the above descriptions is closest to the entropy of the real-world object. That entropy is a theoretical minimum and probably unobservable.

An entropy map would represent, on a scale of increasing entropy, the relative positions of the various conceptual models and the object of interest. The real world object would sit at some theoretical minimum. Each user's model, representing information that is true of the model, sits higher than this minimum. The expert has access to the user's information, an extensive knowledge base, and perhaps data that has not been considered by the users.

The expert's model behaves an improving attempt to satisfy the users' believe and move towards the theoretical minimum. The users take the information being sent from the expert's model and manipulate their own to chase the expert's model downwards. The relative position of the expert's model is controlled by the users; they decide which model is the most likely and what additional information may be needed to resolve a conflict.

These points are better illustrated with the hypothetical problems described earlier in the paper. This time, however, we will considered examples from the expert's, as opposed to the user's, point of view. The first problem is presented by a conventional user to a human expert and is solved. The second problem is presented by an unconventional user to a system which fails to solve it. The reasons for the failure can tell us a lot about both the potential, and the limitations, of expert systems.

4.2 THE USER, AS PERCEIVED BY A HUMAN EXPERT

Back in the scene of Section 2, a user asked a question, "how did fine-grained sandstone get into these cuttings?" To understand how an expert can answer this, it is necessary to go deeper into the demand knowledge of sedimentology.

The age-dependent processes of sedimentation and erosions result in rocks possessing certain associated features. Interpretation in petroleum geology is fortunate in that most of these processes are extremely slow. They are not only slow in terms of the rate of additional or removal of sedimentary material, but they are slow in the sense of the rate of change of the processes themselves. The spatial variation in some of the features of a buried, sedimentary rock can be described from rock cuttings or geophysical logs. The values of these features, and their variations along an interval in the borehold, are particularly indicative of variations in the slow, age-dependent processes which begat them.

Johannes Walther's law of succession of facies is frequently used by geologists to interpret the original environment of sedimentary deposition. (The facies of a rock has a number of definitions. In this paper, it is understood to mean the sum of all the relevant, distinguishing features of a rock.) Commonly, Walther's law is state as "facies sequences observed vertically are also found laterally." We will avoid detailed discussions about the restrictions on applying the law. It does, however, have important implications for relating location-dependent processes with age-dependent processes in sediments.

Consider a sedimentary surface upon which a thickness of sediment, h, is being deposited in an interval of time, t. The rate of deposition is therefore dh / dt. Take a feature value, F, of an object just deposited on this surface. We can translate Walther's law into the terminology of this paper as... (?)

The rate of change of F with thickness, h, in the direction of surface growth is equal to the rate of change of F with time, t, along the path of transportation, divided by the rate of deposition, dh / dt.

As an example, let the rate of change of a sandstone's grain-size feature, going upwards in an old-river bed, be –0.003 mm / foot. Say that, at any one level, and at the time of deposition on the river beg, the rate of position was one foot in a hundred years. The law tells us the rate of change of grain size, for grains being deposited by the river was –0.003 mm / 100 years.

In a non-trivial sense, the supply process is communicating with the growing sedimentary deposit. Some of its features, varying with respect to age of deposit, are going to be preserved in space and can be measured by logging tools like the density device, mentioned in Section. (Of course, the values of the features of a sedimentary rock may be modified by post-depositional processes such as compaction, tectonic stressing , precipitation, chemical and thermal

modifications.) With the necessary knowledge, however, and a record of the variations of a feature in physical space, and expert can make a translation to obtain a picture of the depositional process in time.

Here then, is another mechanism for communication between users. The well logs are continuous records against depth of more-or-less accurate measurements of a combination of rock features. The log analyst provides the expert with opinions on what corrections to the measurements are necessary. The expert also has to know such things as, representations of Archie's equation (indicating the "behaviour" or rock resistivities as grain size changes), the effects that various perturbations might have on the measurements – for example, an error in calibrating the density tool – how the logging tools function, and so on. The sedimentologist teaches the expert about some of the characteristics of deposition in old-river beds. Not just any river bed, but those typical of the region where the well has been drilled. The expert has to learn, in addition, such thinks as typical rates of deposition according to position downstream, and the rate of change of grain size with space and time. The sedimentologist gives him opinion on where, and at what depth, the river beds are likely to be encountered.

The expert takes this mass of information and starts to construct a mental object from user opinions and the measured data. The expert's knowledge is used constantly to keep this mental object geologically consistent. Any translation, or propagation of consequences that contradicts the users' opinions is presented to the user. For example the expert might tell the log analyst that, "approaching the top of this old-river channel, the resistivity log should show an increase and doesn't?" Perhaps the sedimentologist is advised where the top of a river beg ought to be, judging from the resistivity and density curves, and modifies his opinions accordingly. The users keep driving the expert's model around each other's territory, until they can agree where it should be "parked." On the user's agreement, a state description of the model is the expert's best opinion on the right interpretation.

The expert's "best" model is therefore assembles and should contain enough information to explain the presence of "fine-grained sandstone" in a way that suits the questioner. If a system can assemble a model, of the same information content as the human expert, it can exchange messages with the user of equivalent quality.

Of great interest, is what an expert system is probably unable to do. Let us use the incident of the mischievous roughneck to illustrate at least one area where the machine would be incompetent.

4.3 THE USER, AS EXPERIENCE BY AN EXPERT SYSTEM

The roughneck, remember, had added a little seasoning of opal to some cuttings drilled up from an oil well. The discomfiture of the young geologist was obviously the subject of his raucous enjoyment. At the risk of destroying whatever human remains in the situation, let us sit the unlovely pair down for a discussion with an expert system. Assume the geologist has taught the system enough for it to know what minerals to expect and the processes by which they might have originated.

The geologist now asks the system the question, "why is there opal in these cuttings?" The system's glass panel, connected to the geologist's communication system, immediately portrays the bead picture that represent the question. The system starts searching for primitives that can be combined together to reproduce the panel picture in the table blow. The following sequence develops:

a. Opals and cuttings are objects, each possessing a large number of features – a few in common. The words "in these," of the question, implies the objects share the same location feature, i.e. they are at the same place at the time the questions is being asked. Should the system go further and treat the opal as "part of," instead of just "in," the cuttings?

   Opal shares at least one important feature dimension with the cuttings drilled up from a borehole. The mineral percentage of sedimentary rock is always positive and normally distinguishable form the rock's fluid or organic-material content. Opal is basically a pure mineral and rock cuttings contain minerals. Both objects possess a "mineral percentage" so that one could be "part of" the other.

   The subtle distinction in the meaning of a proposition is more important that at first appears. If one object had been alien to the other; for example, if there was an "Aylesbury duckling in these cuttings," a human would look for an alien process that put it there. The human search for a explanation is eased because the human can understand why something is unlikely. The system's search, however, is restricted by what it knows is likely and, in the absence of extensive world knowledge, would have trouble locating unlikely explanations.

b. The system's list of likely minerals, in this particular well, does not include opal so a likely replay might be, "Opal is unlikely." A little "thought," however, would show that this reply brings no useful information to the geologist as it is the reason why he asked the question. Still, as argued before, it does not seem impossible to have an expert system guess at the

> geologist's model. (The system could reasonable model its deficiencies and these deficiencies would not include ignorance of the likely minerals.)
>
> Assuming the system realizes that there is no point in telling the geologist what he probably knows already, it still has the problem of manipulating its models of the "cuttings" object and its impressions of the geologist's model of those "cuttings." The resulting model must accord with the known facts, which include the fact of the geologist's question. One fact the system has not digested is the laughter of the roughneck. It is evident that, if the system could model the roughneck's view of things, it could come up with, "because the roughneck put the opal in the cuttings as a joke." Failing that, the system can only extend its search for a plausible explanation of which the geologist was unaware. If nothing is found, the answer would be, "don't know."

4.4 HUMAN AND SYSTEM EXPERTISE: A COMPARISON

The essential difference between the human and the system expert can be summarized in two possible situations. One, the messages from the human user are explicit and provide all the relevant information about an object; from the expert's point of view, the object can then be modeled equally well by human or system. Two, the messages are implicit, requiring the expert to modify a plausible model, in order to infer the most likely meaning of the message. If, in situation two, the system's model does not have the same dimensions as the user's model, it cannot "understand" the message in the sense transmitted.

Whether this distinction causes significant problems for an expert system, will be a questions of experience and, of the nature of the domain to which it is being applied. Much of human thinking and communicating about geological objects, involves implicit message passing. Fortunately, dimensions, involved in such messages, are the ones used to measure the effects these object have on humans. The dimensions of the user's model are therefore the more-obvious, perceived dimensions of inanimate objects – good reason to believe a system could assemble models compatible with implicit message passing between humans about a restricted domain.

A "funny" situation is an object in the mind of a human; its properties can only be seen by its effects, such as laughter. A system cannot understand a "funny" situation unless it can laugh, experience the effects of that laughter and so, incorporating "funniness" as a dimension in one of its own models. As system could, however, understand that a human finds something funny – it could have been told that explicitly. All those words, involving a peculiarly-human dimensions, will have at least one thing in common – artificial "understanding" system can only treat them as objects of definition. "Funny" would then be

understood by a system as a world that humans use to describe an object at which they laugh. This limitation will apply as long as the system is not equipped with the feature that allow it to represent the human dimensions and treat the worlds as objects of description, e.g. "that object is funny because it makes me laugh."

I have briefly contrasted how the human user of expertise might appear to the human expert and to an expert system. It has been said that the information content of a message depends on its unexpectedness, also that the most efficient messages are almost indistinguishable from the noise of transmission. Humans become extremely efficient at sending each other implicit messages. A single word, a gesture, a slight change in emphasis, or in expression can be understood without conscious effort. It is a basic premise of this paper that the meaning of an implicit message can only be extracted if the receiver can model the state of mind of the transmitter.

A system's ability to understand the human users will be limited in two areas:

a. Its efficiency in assembling a conceptual model that corresponds to a functional image of the user's model.

b. The extent to which it is equipped with the features and associated dimensions that were needed to make up the user's models.

The impact of the ideas, so far discussed, on the design, possibilities, and dangers of expert systems will be discussed in the concluding section.

SECTION 5. CONCLUSIONS

5.1 A THEORY OF EXPERT THINKING: A SUMMARY

Using the domain of petroleum geology to illustrate a number of key concepts, a theory of expert thinking has been developed. For clarity, the major components of human expertise were tentatively identified and then studies as separate issues. In reality, it seems that learning, understanding, effect propagation, and conflict handling are partial descriptions of contemporary mental activities. The mind is doing all these things simultaneously during its attempt to build functional images of the outside world. Success in these attempts can be recognized when, and if, the mind cannot distinguish between the effects upon it of the real world, and of the image it has built of that world. A description of that image, constrained and guided by the questions asked of it, appear to the user as an interpretation of the real world.

The arguments developed, to support this theory of expert thinking, are based on the consequences of certain, unproven propositions. The propositions, and their consequence, are summarized below in their order of appearance in the paper:

a. The building blocks of expert knowledge are primitives made out of features that are attached to behavior patters. In tern, these blocks can be combined into higher level primitives of objects linked by processes. The highest level primitives are called systems. A primitive in this paper is considered to be the smallest mental object to be understood and generates effects having the same information content as the effects used to locate or construct it.

b. The information content of a system's properties is the same as the information content of the system's effects. Properties may be possessed by features, objects or system but only their effects are experienced. The limit on the information that can be extracted from something is equal to the information contained in the sum of the effects it has on its environment.

c. A system understands an external something when it creates an internal object producing the same internal effects as the something to be understood. A consequence of (b) is that, for the understanding system, the external something and the internal object are then isentropic, having the same information content.

d. The problem of how to make a system understand something can be reduced to the problem of how to assemble mental objects having the same internal effects on the system as the thing to be understood. A

number of techniques, useful in manipulating and controlling an appropriate mental-object generator, have been identified:

– Features and behavior patterns may only be associated in a way that preserves dimensional integrity. (“The grain’s texture becomes larger,” “the bed’s azimuth is red, “ are associations that would not be permitted by nature.)

– A behavior pattern normally implies a certain direction to the change in value of a dependent feature with respect to an independent feature. The relative values of the dependent feature must reflect this direction of change. (Whilst being abraded, grains get smaller as time increases.)

– Whatever the independent dimensions of the behavior patterns involved, the value of a dependent dimensions of the same object must have the same value at the same point in time. (If a given rock grain has a small grain size in space, it must have a small grain size in time.)

– Records are made of the feature differences between real-world objects and “similar mental objects as experienced by an understanding system. The differences between real-world and “similar” mental objects as experienced by an understanding system. The differences produced by perturbing a real-world object are reciprocal to the effects of a similar perturbation to its functional image. (This assumes the system has a correct knowledge of a real world that is continuous and deterministic.) Recognition of a set of differences in memory that are reciprocal to those being experienced direct the system to a perturbation that will “probably” restore a condition of “understanding.”

e. A description of an expert’s functional image of a real-world object appears to the users as an interpretation of that object.

f. An expert’s best opinion about an object is the lower entropy description to satisfy the user. The best expert’s best opinion is the one closest to the entropy of the object itself.

g. If all messages about a domain are explicitly, a system can be built to communicate and understand knowledge in the same way as a human. For a system to understand implicit information, the system has to be able to construct mental objects having the same internal effects as the human, sending or receiving the messages. If the system is not equipped to experience the same internal effects, it cannot reach a state of human-like understanding.

5.2 THE DESIGN OF EXPERT SYSTEMS: POSSIBILITIES AND PITFALLS

The ideas described in this paper have numerous implications for the design of expert systems. Every domain that is a potential candidate for such a system occupies some loosely-defined area in a space of total information. Human

being select, or have forced upon them that part of the information space which is considered relevant and interpretable. Just because information is identified as useful does not mean it can be measured, or if measured, that it unambiguously defines that state of the object studied. Domain objects, processes and systems may have been selected by humans for study, but nature determines the information content of the messages available for interpretation.

Expert system design should therefore take into account the nature of the subject matter and the way it broadcasts information about itself. It is a mistake to apply purely deductive techniques to information that can be generated by a large number of system configurations. It is redundant to try and synthesize models of real-world systems if those systems are transmitting enough information to allow an unambiguous interpretation.

Examples of the misapplication of techniques are plentiful; the econometrician's behavioral analysis of economic systems which should be modeled with more dimensions; a medical diagnosis that is based solely on classifying a disease by its symptoms, not by exercising a model of the body to determine the most probable configuration that explains those symptoms; purely deductive attempts to understand natural language when the implicit nature of the messages require a function image of the mental object described, to be assembled in the mind of the receiver.

Properly designed and applied to the appropriate subject, an expert system could provide a power tool for the users to exploit the knowledge peculiar to that subject. Such tools must not appear to present the "truth" when, at best, they can only give a "most probable" association. Instead, they should give the human user the chance to experiment with all that is known, or relevant, about a subject, before making potentially disastrous experiments with reality. An expert system should never take a decision. It should display the consequences of various opinions, to all its different users, whilst keeping the model that performs this function consistent with the sum of the domain knowledge.

The arguments presented in this paper lead me to believe there are no major obstacles to the construction of an expert system providing services at the level of an expert petroleum geologist. If this system can be realized in practice, many of the principle used in its construction will be applicable to those sciences that combine and use knowledge in a way similar to petroleum geology. To temper such optimism, a real concern is that users will not insist upon the give specifications required of a expert system; the capabilities of learning,

understanding, propagating effects, handling conflicts and justifying its results. If they don't insist, they will receive system that are tendentious, frequently wrong and, unfortunately, liable to be uncritically believed.

Through the mouthpiece of his great detective, let us give Sir Arthur Conan Doyle the last world on what a user should insist upon, when considering advice from an expert, human or otherwise. In "A Study in Scarlet," Sherlock Holmes describes, with typical immodesty, his expert skills to Dr. Watson and a respectful audience of policemen:

> *Most people, if you describe a train of events to them will tell you what the result would be. They can put these events together in their minds and argue from them what will come to pass. There are few people, however, who, if you told them a result, would be able to evolve from their own inner consciousness what the steps were which led up to that result.*

David C. Hawkins
Ridgefield, Connecticut USA

July 15, 1981

For further information or to invite David Hawkins for a speech or workshop on the Analysis of Expert Thinking, please contact:

Email: DCM-001@DavidHawkinsResearch.com
Website: www.DavidHawkinsResearch.com

Mail: David Hawkins Research, DCM-001
c/o #202- 2001 East 36th Ave.,
Vancouver, B.C., Canada V5P 1C9

SECTION 6. REFERENCES

1. Sir Arthur Conan Doyle, A Study in Scarlet. New York: Doubleday & Company Inc., Republished.

2. Webster's New Collegiate Dictionary. Springfield: G&C Marriam Company, 1977.

3. Glossary of Geology. Washington: American Geological Institute, 1973.

4. Harvey Blatt, Gerard Middleton, Raymond Murray, Origin of Sedimentary Rocks. Englewood Cliffs: Prentice Hall, Inc., 1972.

SECTION 7. ACKNOWLEDGMENT

Many of the ideas in this paper have developed from discussions with Luc Steels, during our joint development of ANALOG, an expert system in petroleum geology that is currently under development in Schlumberger-Doll Research, Connecticut. For all his suggestions and critical advice, I would like to express my thanks. Ray Campbell, Darwin Ellis, and Luc Steels have read the paper and I am grateful for their recommendations of changes, most of which have been incorporated.

The original source for my theories of expert thinking has been observations of real-world experts, in action on a range of problems of different domains. I would like to think them collectively, but anonymously, for revealing some of their techniques – they allows me to construct a hopefully plausible and implementable theory of what experts do when they provide users with the results of their expertise. If the theory proves to be a poor representation of expert thinking, I hope it can still achieve two goals: expert systems users will be encouraged to better define their requirements and, expert-system designers will become aware of the considerable risks of translating the vernacular into the oracular.

For further information or to invite David Hawkins for a speech or workshop on the Analysis of Expert Thinking, please contact:

Email: DCM-001@DavidHawkinsResearch.com
Website: www.DavidHawkinsResearch.com

Mail: David Hawkins Research, DCM-001
c/o #202- 2001 East 36th Ave.,
Vancouver, B.C., Canada V5P 1C9